Visualizing multifractal scaling behavior: A simple coloring heuristic

A. C. Gilbert W. Willinger A. Feldmann
AT&T Labs–Research
180 Park Avenue
Florham Park, NJ 07932
Email: {agilbert,walter,anja}@research.att.com

Abstract

The self-similar or monofractal scaling behavior of a signal can be readily visualized in the time domain by plotting and comparing the signal at many different levels of resolution (i.e., time scales). No comparable visualization techniques appear to be available for qualitatively assessing the presence of multifractal scaling in a given signal. To help in developing an intuitive understanding of multifractals, we propose in this paper a simple coloring heuristic that gives rise to an effective, time domain-based visualization technique for multifractals. Our approach combines wavelet-based statistical denoising techniques with wavelet characterizations of local regularity for functions. We illustrate our heuristic with a number of examples, including the binomial measure, a simple and well-studied multifractal, and measured traces of actual data network traffic.

1. Introduction

Multifractals provide a mathematical framework for detecting and identifying complex local structure, for describing local singularities (e.g., see Evertsz and Mandelbrot [6]). With time-dependent scaling laws, they are more flexible in describing locally irregular phenomena than monofractals, which are governed by single time-independent scaling laws. Exactly self-similar processes are special cases of monofractals; their degree of local irregularity is the same across all scales and can be captured by a single global parameter, the *Hurst parameter* H. Multifractal structures have been found in a wide variety of physical systems from turbulence [2] to rain and rain clouds [9, 15], and more recently, Internet and other data network traffic [21, 16, 7].

Statistically, monofractals or exactly self-similar processes exhibit identical characteristics at all (or over a wide range of) time scales. This property gives them the appearance of "looking the same across many different scales," which in turn offers ample opportunities for simple and intuitively appealing visualization heuristics for monofractal scaling behavior in the time-domain (see for example Figure 4 in [14]). For multifractals, such time domain-based visualization heuristics are either non-existent or are rarely used in practice. At the same time, there exist many different mathematical methods for describing multifractal scaling behavior in a signal, but they are, in general, difficult to interpret and delicate to calculate. For example, one of the most commonly used methods is based on the *multifractal spectrum*, a global summary statistics (i.e., histogram or "spectrum") that gives the frequency with which different local scaling exponents occur in the underlying signal. However, because the multifractal spectrum is defined geometrically in terms of the Hausdorff dimension of the sets where the signal has a common scaling behavior, estimating it from measured data is usually difficult. In practice, it is approximated and estimated via the Legendre transform of the *structure function*, where the latter captures the scaling behavior of the higher moments of the data. Arneodo et al. [1], Jaffard [12, 13], Riedi [18, 20], and Gilbert et al. [8] introduce a number of different methods for constructing the structure function of singular signals and measures, including several wavelet-based methods. However, while the main objective of all these methods is to provide a mathematical description of the multifractal scaling behavior of a signal, none of them results in a simple and intuitive heuristic for visually distinguishing monofractal from multifractal signals in the time domain.

In this paper, we present a heuristic coloring method that can be naturally incorporated into conventional time series plots and that describes visually multifractal scaling or local irregularities in a signal. We show how to color the graph of a function (e.g., time series) differently in the regions where it is smooth from the regions where it is non-smooth or irregular. More precisely, the colors reflect the gradations in the (estimated) local Hölder exponent of the function and give rise to a simple time series plot that incorporates information about the degree of local irregularities in the signal. The range of colors shows the range of local irregularity and the distribution of the colors within the signal illustrates the distribution in time of the local irregularities. Our coloring heuristic combines the local regularity characterization properties of wavelets with the statistical (wavelet-based) techniques of denoising. To this end, Sections 2 and 3 outline the mathematical foundation of wavelet characterizations of local irregularity for functions and illustrate the basic idea behind our approach with a simple example. A brief discussion of the statistical theory of denoising is given in Section 4. In Section 5, we put the pieces together and present and illustrate the heuristic coloring method. We conclude in Section 6 with some observations about the nature of the technical results that are motivated by our simple coloring heuristic and would put the proposed approach on a solid mathematical basis.

2. Local scaling and wavelet coefficients

We base our heuristic method for visualizing local scaling on the wavelet transform. Because this transform divides data into different frequency components and analyzes each component with a resolution matched to its scale, we can use the coefficients of a wavelet decomposition to collect microscopic information about the scale-dependent properties of the data. The wavelet transform can scan the singularity structure of both functions and singular measures. Arneodo et al. [1] (see also the references within) apply wavelet-based techniques to the study of multifractals and turbulence and Jaffard [12, 13] make rigorous a number of the wavelet-based multifractal "tools." Gilbert et al. [8] use a related wavelet-based tool to investigate the scaling structure of conservative cascades, a particular class of random multifractal measures, and they apply these methods to explore the multifractal nature of data network traffic. In this section we discuss the mathematical basis of the wavelet transform as a microscope. We illustrate the algorithms that these theorems suggest, and we briefly explore their use in practice.

2.1. Mathematical foundation

One of the underlying principles of time-frequency decomposition is that if a signal X is locally regular about a point t_0 in time, its wavelet coefficients near t_0 decrease with scale. To formalize this tenet, we say that X is *Hölder continuous with exponent* α, $0 < \alpha < 1$, at t_0 if as $h \to 0$,

$$|X(t_0 + h) - X(t_0)| \leq C|h|^\alpha.$$

Let us assume that the wavelet $\psi \in C^1$ and that ψ has compact support. Then the decay character of the discrete wavelet transform of X reflects the Hölder exponent of X at t_0 (see [11, 4, 17, 10] for a more detailed discussion):

Theorem 2.1 (Jaffard [11]) *If X is Hölder continuous with exponent α, $0 < \alpha < 1$, at t_0, then*

$$|\langle X, \psi_{j,k}\rangle| \leq C 2^{-j(\alpha+1/2)}(1 + |t_0 - 2^{-j}k|)^\alpha \qquad (1)$$

for all j. Conversely, if (1) holds and if X is known to be C^ϵ for some $\epsilon > 0$, then

$$|X(t_0 + h) - X(t_0)| \leq C|h|^\alpha \log \frac{2}{|h|}. \qquad (2)$$

Unfortunately, we do not have exact equivalence between (1) and (2). In fact, the estimate (2) is an optimal estimate. If we omit the condition $X \in C^\epsilon$ and assume instead that X is only continuous or if we omit the logarithm in the above estimate, we can find counterexamples. If we modify the relation (1), we have the following result:

Theorem 2.2 (Jaffard [11]) *Define for $\epsilon > 0$, the set*

$$S(t_0, j; \epsilon) = \Big\{ k \in \mathbf{Z} \,\Big|\, \mathrm{supp}(\psi_{j,k}) \cap (t_0 - \epsilon, t_0 + \epsilon) \neq 0 \Big\}.$$

If for some $\epsilon > 0$ and some α, $0 < \alpha < 1$,

$$\max_{k \in S(t_0,j;\epsilon)} \Big|\langle X, \psi_{j,k}\rangle\Big| = O(2^{-j(\alpha+1/2)}), \qquad (3)$$

then X is Hölder continuous with exponent α in t_0.

Similar results can be shown for Hölder exponents $\alpha > 1$, including for $\alpha \in \mathbf{N}$ (using the Zygmund classes). One may also prove similar theorems if the wavelet ψ does not have compact support, provided ψ and ψ' decay quickly at infinity. In addition, similar results hold for the continuous wavelet transform of X. See [17, 10] for more technical and detailed discussions.

We may study the singularities of a measure μ in a fashion analogous to the local Hölder exponents of a function. We measure the singularity exponent $\alpha(t_0)$ of μ at the point t_0 via

$$\lim_{l \to \infty} \frac{\log\big(\mu(B_{2^{-l}}(t_0))\big)}{-l \log 2} = \alpha(t_0).$$

where $B_{2^{-l}}(t_0)$ are the dyadic intervals of size 2^{-l} that contain the point t_0. If the above limit does not exist, we leave $\alpha(t_0)$ undefined. We say that μ is *multifractal* if the scaling behavior of μ at t_0 depends on the point t_0; that is, if $\alpha(t_0)$ varies as t_0 varies. A measure μ is *monofractal* if a single global scaling exponent α describes the singularities of μ. Analogous theorems hold for the (continuous) wavelet transform of a measure μ; if μ has local scaling exponent $\alpha(t_0)$ at t_0, then the (continuous) wavelet transform of μ at t_0 also has $\alpha(t_0)$ scaling behavior (see [1]).

2.2. Practical implications

To illustrate the above theorems and to show how we use the discrete wavelet coefficients to characterize the local regularity of functions, we analyze the periodic function given by

$$X(t) = \begin{cases} 32\sqrt{-t}, & \text{if } -1 < t < 0, \\ 32\sqrt{t}, & \text{if } 0 \leq t \leq 1, \end{cases}$$

and shown in Figure 1 (top). This function X is a piecewise smooth function with two distinguished points; one singularity (at $t = 0$) is a cusp discontinuity with Hölder exponent 1, and the second singularity (at $t = -1 = 1$) has Hölder exponent between 1 and 2 since the derivative of X at $t = 1 = -1$ does not exist but it does not blow up as at the cusp. The time/scale decomposition of $X(t)$ is shown in Figure 1 (bottom). The brightly colored “fingers” in the time/scale plot “point” to the singularities in the function, making it easy for us to “see” where the function is irregular and how irregular it is in the different points in time. With this artificial example, we can easily pick out the discontinuities from the graph of the function itself. The brightest finger (and the slowest decaying wavelet coefficients) is at the cusp discontinuity. Note that we have used periodic boundary conditions. The finger pointing to the cusp singularity is brighter than those observed at the edges. Also, the decay of the wavelet coefficients above smooth points in X is much faster than above the singularities.

To characterize the local regularity of the function using the (discrete) wavelet coefficients, we use Theorem 2.2 and plot $\log_2 A_j(t_0)$ as a function of scale j, where $A_j(t_0) = \max\{|\langle X, \psi_{j,k}\rangle| \, | t_0 \in \mathrm{supp}(\psi_{j,k})\}$, and calculate the slope of $\log_2 A_j(t_0)$. Unfortunately, this method is highly dependent upon the location of the point t_0; if t_0 is not a dyadic rational, then we need to take many fine scales of wavelet coefficients in order to fit a line through $\log_2 A_j(t_0)$. In other words, it may take many fine scales before the (discrete) wavelet coefficients decay in a regular fashion. See [4] for a discussion of this phenomenon. If we use a continuous transform, we avoid this dependence upon location. For numerical purposes, we compromise and use a redundant translation-invariant discrete wavelet transform; we use a non-decimated discrete wavelet transform which is oversampled at coarse scales. Again, we plot the logarithm of the magnitude of the (redundant) wavelet coefficient above a point t_0 as a function of scale and calculate the slope of this function. Figure 2 shows this function for a smooth point of X to the right of the cusp singularity. Notice that we still need to discount the first several coarse scale coefficients to calculate the local regularity for the smooth point.

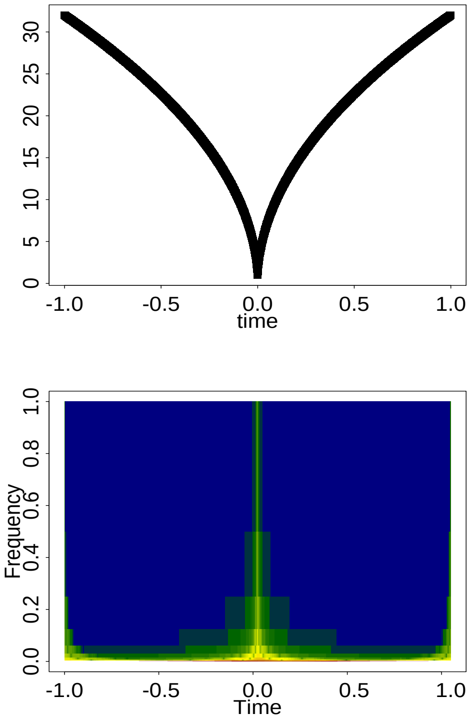

Figure 1. Test function X (top) and its time-scale decomposition (bottom).

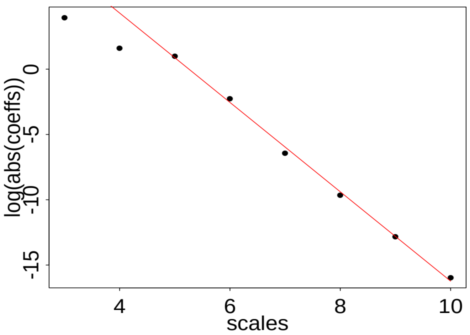

Figure 2. Decay of wavelet coefficients above a smooth point to the right of the cusp singularity.

3. Coloring according to local regularity

In the previous section we showed how to estimate the local regularity of a function (or a measure). Next we illustrate how to use this information to visualize the range of regularity in a function or to compare visually the regularity of two signals.

Using the same simple function X as above, we begin with the non-decimated wavelet transform of X, and at each point we use the previous discussion to estimate the scaling exponent corresponding to that point. Next we divide the scaling exponents into quantiles and color each point according to what quantile its scaling exponent belongs. We use a gradation of color, the darkest color for the lowest quantile and the lightest color for the top quantile. We plot in Figure 3 the signal X colored according to its local regularity, the estimated scaling exponent at each point in X (middle), and the density of scaling exponents less than $\alpha = 3$ found in X (bottom). Because we used the "s8" wavelet for the non-decimated wavelet transform of the function X, and because "S8" has three vanishing moments, support width 7, one derivative, and is almost symmetric, we find that the smooth pieces of X have scaling exponent $\alpha \approx 3$. However, we know that those pieces are smooth and our estimate is limited by the number of vanishing moments of the analyzing wavelet. In other words, we can say that the smooth pieces of X have scaling exponents at least $\alpha = 3$. For this reason, we plot the density of scaling exponents α that are no more than 3. Since X is smooth everywhere but at two points, almost all of the scaling exponents are equal to 3. There are a few tiny bumps in the histogram of scaling exponents for $1 \leq \alpha \leq 2$ and the middle plot in Figure 3 shows there are indeed several distinguished points where $1 \leq \alpha \leq 2$ while the majority of the scaling exponents are around $\alpha = 3$.

This procedure suggests a class of coloring algorithms, consisting of the following five main steps:

- compute redundant (non-decimated) wavelet transform,
- estimate scaling exponent at each point,
- divide all scaling exponent estimates into quantiles,
- color each point according to what quantile its scaling exponent belongs,
- generate a histogram (or density plot) of all the scaling exponent estimates.

4. "Denoising" data

While the coloring procedure works very well for simple functions which are essentially smooth with a few exceptional points, this procedure fails on more complicated signals, including multifractal measures and actual data. This failure is mainly due to the fact that it is generally very hard to measure the decay rate of the wavelet coefficients above a particular point in bursty signals—the wavelets "feel" the

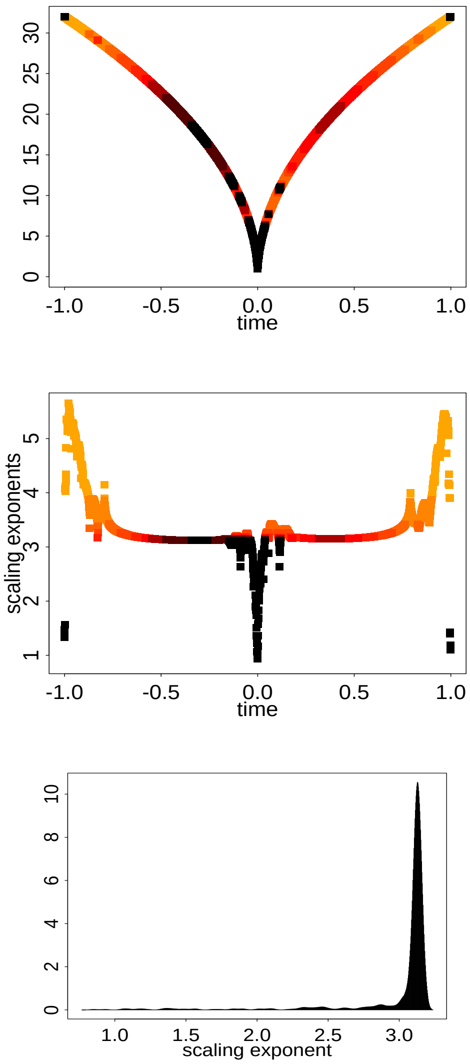

Figure 3. Test function X colored according to local regularity (top), estimated scaling exponent at each point in X (middle), and the density of scaling exponents(bottom)

bursts of the neighboring points. To circumvent this problem, we "denoise" the bursty signal before estimating its scaling exponents, removing all but the largest spikes, creating a signal that is more like the simple piecewise smooth function X. We must emphasize that the pieces of the bursty signal we remove are not noise in the traditional sense!

There are several important reasons for "denoising" signals. One reason is statistical prediction or function estimation and the dual is compression (whether for efficient storage and transmission or efficient numerical algorithms).

Denoising is used for estimating an unknown function X from noisy data Y where $Y(t) = X(t) + N(t)$ and the noise $N(t)$ is Gaussian with variance σ^2, although there are many generalizations of this simple signal-plus-noise model. The dual of denoising for estimation is nonlinear approximation or compression where one throws away insignificant wavelet coefficients and keeps only the significant ones. The essence of wavelet "denoising" lies in the ability to represent locally nonsmooth phenomena with a few significant wavelet coefficients.

Donoho and Johnstone's WaveShrink algorithm (see for example [5, 3]) consists of three steps:

- compute (discrete) wavelet transform of data,
- shrink wavelet coefficients to zero,
- invert wavelet transform.

This algorithm works very well if the original signal has a sparse wavelet representation, and it has a large number of uses and generalizations. It relies on three basic premises: (i) the original signal features are represented by a few wavelet coefficients (hence wavelets are good building blocks for the signal), (ii) noise affects all wavelet coefficients, and (iii) shrinking the coefficients to zero removes the noise while preserving the features of the signal.

With the signal-plus-noise model $Y(t) = X(t) + N(t)$, the wavelet transform, $WY = \{d_{j,k} | j = 0, \ldots, J-1\}$, of the data is the sum of the wavelet transforms of the original signal and of the noise; i.e., $WY = WX + WN$. In addition, the wavelet coefficients of the noise are independent Gaussian random variables with variance σ^2. We shrink the data wavelet coefficients to zero at the j finest scales, by setting $\widetilde{d}_{j,k} = \delta_{\lambda_j \sigma_j}(d_{j,k})$. The threshold λ_j and the estimate of the scale of the noise σ_j are shrinkage parameters. There are two basic shrinkage rules, the *soft-thresholding* and the *hard-thresholding*. The hard-thresholding shrinks to zero all coefficients $d_{j,k}$ which are less than λ_j in magnitude; the soft-thresholding is continuous because it shrinks coefficients above the threshold value λ_j (i.e., it "rescales" them). More specifically, the shrinkage rules for hard- and soft-thresholding (respectively) are given by

$$\delta_{\lambda\sigma}(x) = \begin{cases} 0 & \text{if } |x| \le \lambda, \\ x & \text{if } |x| > \lambda \end{cases} \qquad \text{and}$$

$$\delta_{\lambda\sigma}(x) = \begin{cases} 0 & \text{if } |x| \le \lambda, \\ \mathrm{sign}(x)(|x| - \lambda) & \text{if } |x| > \lambda\sigma. \end{cases}$$

The default options for the algorithm are to shrink the $j = J - 2$ finest levels with soft-thresholding using (i) a uniform threshold $\lambda_j = \sqrt{2 \log n}$ (where n is the number of data points) and (ii) an estimate of the scale of the noise $\sigma_j = \hat{\sigma}(d_J)$ that is based on the variance of the finest level coefficients.

Figure 4 shows two denoised versions of the deterministic binomial measure (binomial measure, for short), one of the best-known and simplest multifractals (e.g., see[19] for a thorough introduction to the binomial measure). The tenth stage of the binomial measure is depicted in blue and was computed with weights $p_0 = 0.4$ and $p_1 = 0.6$. The denoised version of the binomial measure using hard thresholding and a large universal threshold value (based upon the $(1-1/n)$th quantile of the coefficients) is shown in purple, and we use soft thresholding with a small (adaptive) threshold value to approximate the binomial measure in red. Note that while the soft thresholding with an adaptive value generates a signal that is very close to the original multifractal, the hard thresholding where we retain only a very small number of coefficients coefficients is very smooth and captures only a few of the spikes in the original signal. We want to again emphasize here a philosophical distinction between our use of the denoising algorithm and its original purpose; our data are not noisy, each point is genuine and is part of the original signal. We threshold the signal or throw away insignificant wavelet coefficients to get a better grasp on the salient features of the signal, to estimate the scaling exponents of those remaining, strong features.

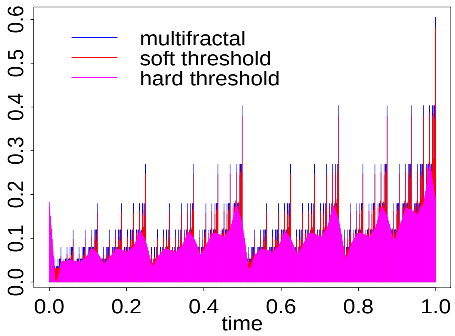

Figure 4. The binomial measure (blue) and two denoised versions—soft thresholding with adaptive threshold (red) and hard thresholding with high threshold (purple).

5. Putting the pieces together

By combining the above denoising algorithm with the intuitive coloring method presented earlier in Section 3, we obtain a simple scaling heuristic for visualizing the local regularity of such bursty signals as the multifractal binomial measure and network traffic traces in the time domain.

To color each point in the signal and to generate the corresponding histogram of the scaling exponents (that is, obtain an estimate of the multifractal spectrum), we use the following modified version of the coloring heuristics:

- denoise data using the basic WaveShrink algorithm,
- color signal based on the denoised data,
- generate histogram of scaling exponents.

To illustrate our method for visualizing multifractal scaling behavior, we use the default settings for the WaveShrink algorithm to color the binomial measure, a genuine multifractal object with known multifractal spectrum (see below). We show in Figure 5 the effect of using no thresholding (top), relying on hard thresholding (middle), and applying soft thresholding (bottom) for this signal. Also, in Figure 6, we compare the histograms of scaling exponents for the denoised versions of the measure and the original measure. Notice that the soft thresholding produces a signal that looks very close to the original measure both visually (i.e., in the time domain) and in terms of the histogram of estimated scaling exponents, while the hard thresholding produces a signal that looks visually (focus here is on color) only slightly like the original signal and has a different histogram of scaling exponents. It is surprising that the main difference in the histogram of estimated scaling exponents is the placement of the center of mass rather than the general shape of the distributions. Figure 6 also shows the known (theoretical) multifractal spectrum of the binomial measure (in black), a parabola with its maximum at one and its extremal Hölder exponents at $\alpha = -\log(p)/\log 2$ and $\alpha = -\log(1-p)/\log 2$, where $p \in (0, 1/2) = 0.4$ is the weight used to generate the multifractal. This plot demonstrates the potential of providing a solid mathematical framework for turning our simple scaling heuristics into a statistically rigorous method for estimating the multifractal spectrum of a given signal. Note that the estimated spectrum does not coincide exactly with the theoretical spectrum in part because we have a limited number of stages (ten) in the multifractal.

Finally, to demonstrate the effectiveness of the proposed scaling heuristics for visualizing and comparing multifractal and monofractal scaling behavior in actual data, we use in Figure 7 two different types of measured traffic traces. The top two plots show a 10-second long segments of measured traffic rates (i.e., number of packets per 10 milliseconds) in a local area network (top) and in a wide area network (second), together with their denoised versions (using soft thresholding). For more details about these traces and the particular networking environments where they were collected, see [14, 7]. The two bottom plots show the results of applying our coloring heuristic. Using an exactly self-similar trace as benchmark (not shown here), we

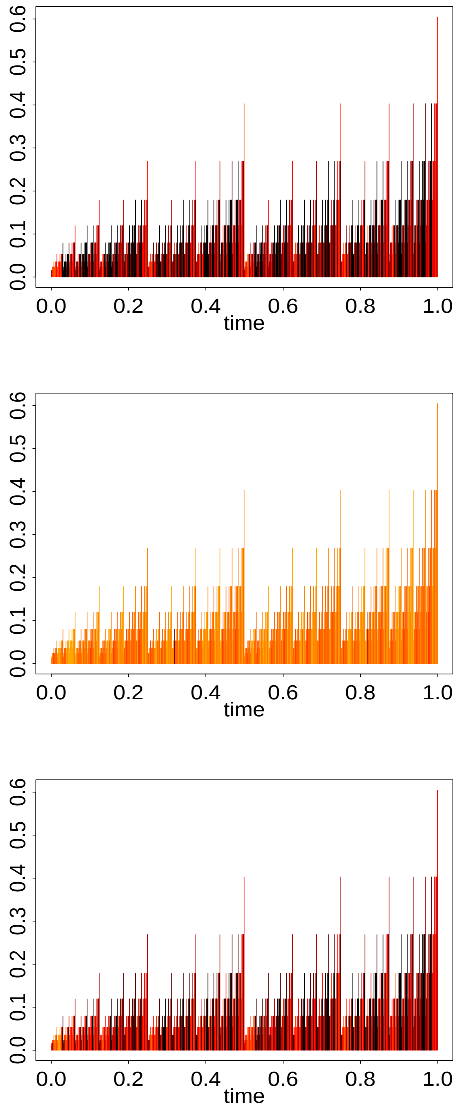

Figure 5. Multifractal binomial, colored without denoising first (top), colored with hard thresholding high value (middle), colored with soft thresholding and adaptive value (bottom)

can conclude that the local area trace is consistent with monofractal scaling and that the observed gradations in color are part of the natural variability associated with estimating local scaling behavior from actual data. In contrast, the visual appearance of the trace in the bottom right plot strongly suggests the presence of multifractal scaling in the wide area trace at hand. Both observations are consistent with the results obtained from a multifractal analysis of

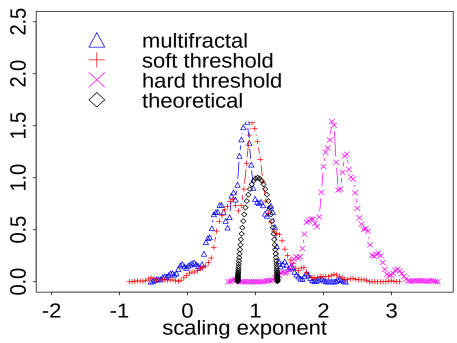

Figure 6. Multifractal binomial: the distribution of scaling exponents from each denoising scheme and the theoretical multifractal spectrum

these traces and indicate the practical use of our proposed method for qualitatively assessing multifractal scaling in a given signal.

6. Conclusion

To develop an intuitive understanding of multifractals, to capture some their qualitative features, and to visualize local irregularity, we introduced in this paper a simple coloring heuristic. This heuristic combines the mathematical theories of statistical denoising and wavelet characterization of local irregularity with visualization techniques to provide a new way of describing multifractal scaling behavior and local irregularity in the time domain. In fact, our approach can be viewed as a visual counterpart for multifractals to visualizing self-similar scaling via the intuitive observation that "monofractals look the same across a wide range of scales."

While the coloring heuristics presented in this paper appear to yield simple and ad hoc visualization techniques, they do give rise to some interesting theoretical problems associated with estimating the multifractal spectrum. The steps in our heuristic approach suggest a way for computing theoretically the multifractal spectrum of a signal that is different from the conventional structure function approach. Indeed, the proposed method and some preliminary experimental results indicate that by calculating the histograms of scaling exponents via the wavelet transform of the denoised signal along a sequence of threshold values that tends to zero, one ought to recover the multifractal spectrum in the limit. Such a result is theoretically plausible for there is a variety of results about the wavelet characterization of local regularity and the pointwise convergence of the denoising schemes has been recently proved by Tao [22]. In summary, the coloring heuristic presented in this paper for visualizing multifractal scaling behavior in a signal is not only simple and intuitively appealing but has also the potential for leading to mathematically new and interesting findings.

References

[1] A. Arneodo. *Wavelets: Theory and Application*, Wavelet analysis of fractals: From the mathematical concept to experimental reality, pages 349–502. Oxford University Press, New York, 1996.

[2] R. Benzi, G. Paladin, G. Parisi, and A. Vulpiani. On the multifractal nature of fully developed turbulence and chaotic systems. *J. Phys. A*, 17:3521–3531, 1984.

[3] A. Bruce and H.-Y. Gao. *Applied Wavelet Analysis with S-Plus*. Springer-Verlag, New York, 1996.

[4] I. Daubechies. *Ten lectures on wavelets*. SIAM, Philadelphia, 1992.

[5] D. Donoho and I. Johnstone. Ideal spatial adaptation via wavelet shrinkage. *Biometrika*, 81:425–455, 1994.

[6] C. J. G. Evertsz and B. B. Mandelbrot. *Chaos and Fractals: New Frontiers in Science*, Multifractal measures. Springer-Verlag, 1992.

[7] A. Feldmann, A. C. Gilbert, and W. Willinger. Data networks as cascades: Investigating the multifractal nature of Internet WAN traffic. In *Proc. of the ACM/SIGCOMM'98*, pages 25–38, Vancouver, B.C., 1998.

[8] A. C. Gilbert, W. Willinger, and A. Feldmann. Scaling analysis of conservative cascades, with applications to network traffic. submitted, 1998.

[9] V. K. Gupta and E. C. Waymire. A statistical analysis of mesoscale rainfall as a random cascade. *Journal of Applied Meteorology*, 32:251–267, 1993.

[10] M. Holschneider. *Wavelets: An Analysis Tool*. Clarendon Press, Oxford, 1995.

[11] S. Jaffard. Exposants de Hölder en des points donnés et coéfficients d'ondelettes. *C. R. Acad. Sci. Paris*, 308:79–81, 1989.

[12] S. Jaffard. Multifractal Formalism for Functions Part I: Results Valid for All Functions. *SIAM J. Math. Anal.*, 28(4):944–970, 1997.

[13] S. Jaffard. Multifractal Formalism for Functions Part II: Self-similar Functions. *SIAM J. Math. Anal.*, 28(4):971–998, 1997.

[14] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the self-similar nature of ethernet traffic (extended version). *IEEE/ACM Transactions on Networking*, 2:1–15, 1994.

[15] S. Lovejoy and D. Schertzer. Multifractals, universality classes and satellite and radar measurements of cloud and rain fields. *J. Geophys. Res.*, 95:2021–2034, 1990.

[16] P. Mannersalo and I. Norros. Multifractal analysis of real atm traffic: A first look. Technical Report COST257TD, VTT Information Technology, 1997.

[17] Y. Meyer. *Wavelets and Operators*. Cambridge University Press, 1992.

[18] R. Riedi. An Improved Multifractal Formalism and Self-similar Measures. *J. Math. Anal. Appl.*, 189:462–490, 1995.

[19] R. Riedi. Introduction to multifractals. Preprint, 1997.

[20] R. Riedi, M. Crouse, V. Ribeiro, and R. Baraniuk. A Multifractal Wavelet Model with Application to TCP Network Traffic. preprint, 1998.

[21] R. H. Riedi and J. Levy-Vehel. Tcp traffic is multifractal: A numerical study. Preprint, 1997.

[22] T. Tao. On the almost-everywhere convergence of wavelet summation methods. *Applied and Comp. Harm. Analysis*, 3:384–387, 1996.

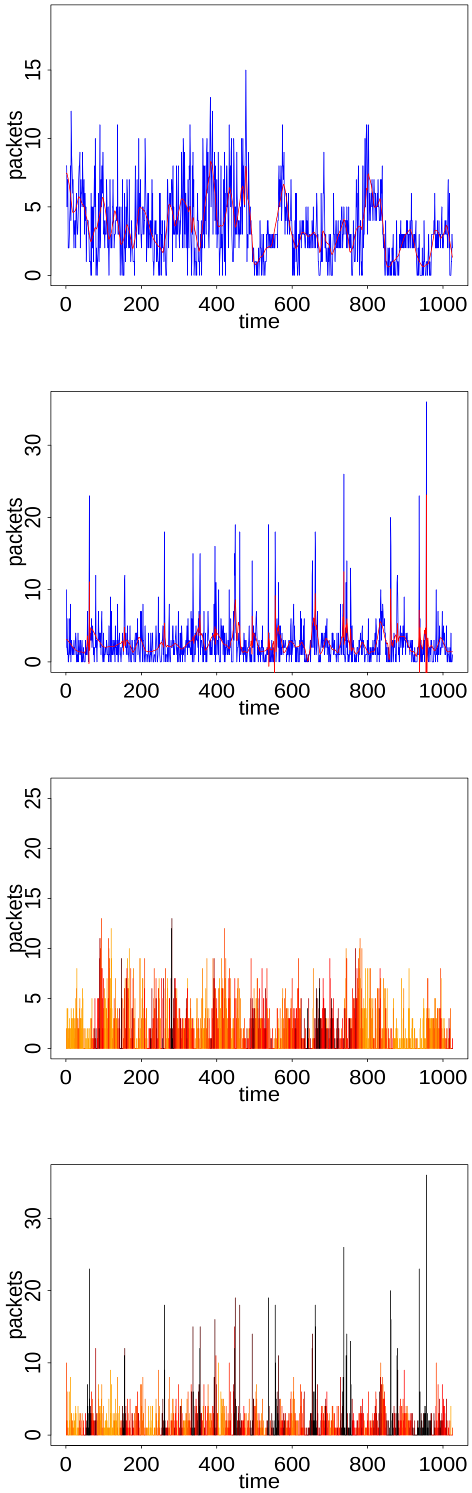

Figure 7. LAN trace denoised (top) and colored (third), WAN trace denoised (second) and colored (bottom).