WHAT IF THE NOISE IS THE PATTERN?

David B. Adams
National Life Insurance Company of Vermont

Fractal geometry has gained widespread acceptance in modeling complex subjects ranging from ecosystems to economics. Harry Erwin (1988) has suggested thut related concepts could be of value for computer measurement. The goal of this paper is to provide an example application of fractal geometry to capacity planning. The example is intended to pique interest and provoke more study in this promising area. The example uses a fractal technique called rescaled range (RIS) analysis, applied to actual workload data. Analysis draws very heavily on work done by Edgar Peters (I 991, 1994) bused on pioneering efforts by Hurst (1951) and Mandelbrot (I 977).

INTRODUCTION

It is surprising that workload study and fractal geometry have not been paired before now. Frdctal geometry is well suited to characterizing what is commonly called "noisy data." Workload data is notoriously noisy, and is usually subjected to extensive smoothing to make it manageable. Figure 1 shows a typically noisy plot of workload activity over time. Activity shown is based on transaction counts for ordinary insurance industry business activities, like requests for customer information or entry of policy data for a new customer.

Obviously. it is not easy to pick out a pattern that describes a workload from such noisy data. Smoothing reduces extraneous roughness, or noise, and reveals an underlying pattern, but an important question arises: what if the noise IS the pattern?

This question is important because it challenges the ideal pursued by conventional methods. Ideally, smoothing techniques result in a straight line: a simple object with a simple, non-fractional Euclidean dimension of one. A jagged, noisy line is more than one dimensional, but less than two dimensional. It has something in between, which Mandelbrot ([MAND77], pp. 14-15) calls a fractal dimension. If the noise actually is the pattern, then smoothing techniques iron important characteristics out of the data.

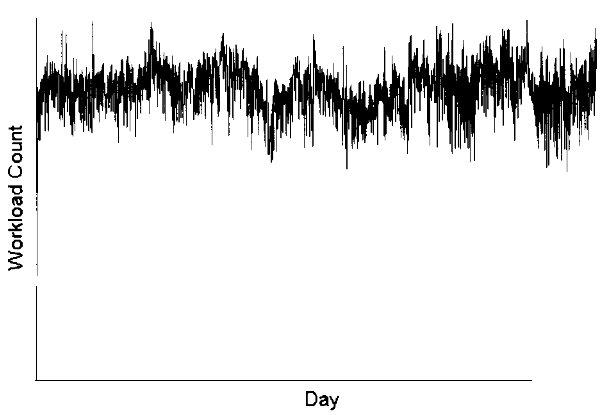

Figure 1. XY plot of workload

It is this fractal dimension that leads to the idea that workload study and fractal geometry are well suited. The relatively new branch of mathematics which is designed to deal with fractal dimensions is fractal geometry.

A FRACTAL TOOL

The fractal tool selected for this paper is called rescaled range (R/S) analysis. R/S analysis, developed by Hurst [HURSSI] to study reservoir capacities, provides a means of describing centrality and dispersion in a time series with one statistic. Description of the R/S analysis process is the main focus here, although attempts have been made at explaining some of the significance of the statistic.

To assist in attaching meaning to the R/S statistic. conventional workload characteristics have been provided. These statistics provide a reference point for comparisons between different approaches.

It should be noted at the outset that although interpretations of R/S analysis are offered, lack of prior application of fractal techniques to workload characterization limits the scope of the interpretations. The conclusion of this paper contains observations regarding alternative fractal approaches and R/S related results, issues, and recommendations.

All discussion is based on the workload data pictured in Figure 1. This data represents daily totals of transactions for regular weekdays over approximately five years.

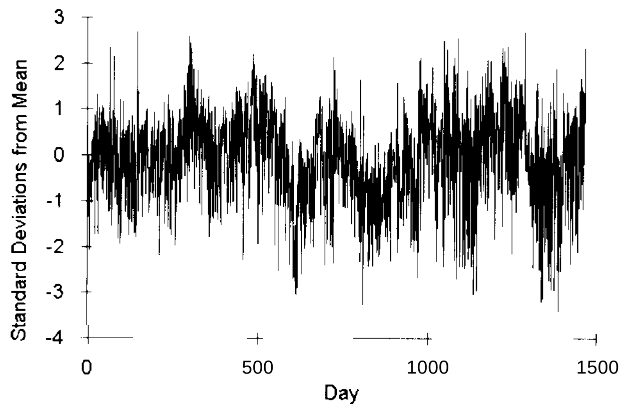

Figure 2. z-scaled workload

TRADITIONAL CHARACTERIZATION

Numerical representation of workload is based on processes which smooth noise out of the data by which the workload ix characterized. For the sake of simplicity. this paper deals with a subset of the rich array of traditional characterization tools. Selected characteristics include simple trending via linear regression. arithmetic mean. and standard deviation (for a useful overview of' these technique\ see [FRIE88m]).

One other statistical technique figures heavily in this paper. Comparison of dissimilar magnitudes is useful here. The description of magnitudes is not being questioned: description of variability is. The comparison tool. called z-scaling. scales the data in relation to its own mean and standard deviation. The result preserves the variability of the original data. but removes the original magnitudes. Figure 3 is a plot of the real workload data from Figure 1 after z-scaling.

Another important side benefit of z-scaling is that results can be compared numerically while respecting the confidentiality of the specific volumes in the original data. No matter what the original values are, the z-scaled data have a standard deviation of 1 and a mean of zero. Only people knowing the standard deviation and mean of the original data can reconstruct it from the z-scaled version.

Traditional characteristics of the z-scaled data appear in Figures 3 and 4. In Figure 3 the heavy. solid line shows the distribution of the data around the mean (which lies along the regression line in Figure 4). The shape of the plot is reasonably close to the classic "bell curve" of a normal distribution, which is shown as a light. dotted line. This suggests that the mean is fairly representative of the data.

Differences between the normal distribution and the data, however, lead to some concerns. In a true normal distribution over 99.73% of the values are within three standard deviations of the mean. Figure 3 shows that the data has a higher than normal percentage of values less than three standard deviations below the mean. Lower than expected values give cause to examine events at the affected times to explain the exceptions.

In Figure 4 a regression line has been drawn for the z-scaled data. The line has a slope of zero, and is just distinguishable along the mean at zero on the vertical axis. The slope of zero indicates that the real workload with which we began in Figure 1 has not changed much in the time illustrated.

Visual examination causes some skepticism about the preceding workload characterization. In spite of the extreme noise in the plot. Figure 4 simply does not look lihe truly random. normally distributed data.

Figure 5 is a plot of the original data scrambled into random order. This rearranged data does have the appearance expected from random data. Data pointx appear randomly, but uniformly, dispersed across the entire time series.

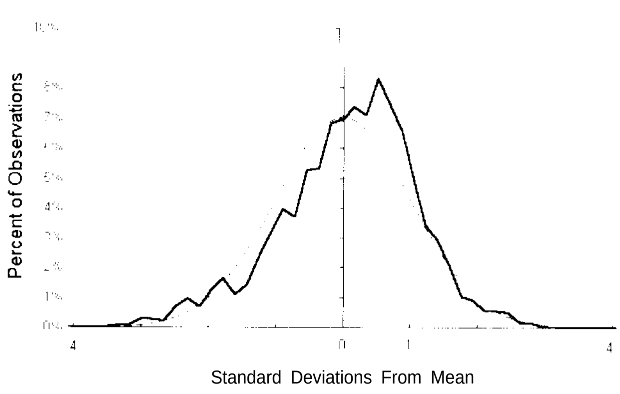

Figure 3. Distribution of z-scaled workload

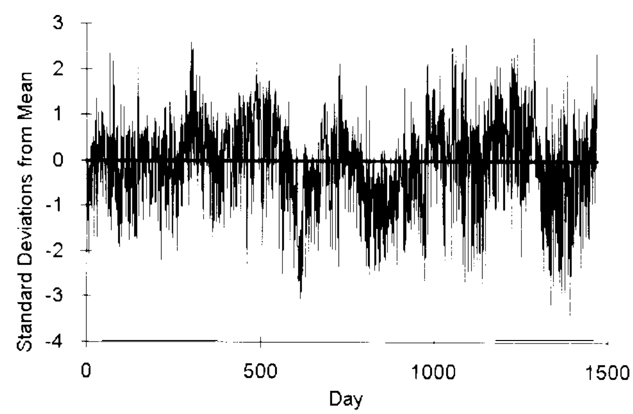

Figure 4. z-scaled data with regression line

It is important to note that, even though Figures 4 and 5 are visibly different from each other, they are still the same data. They share the same distribution. mean. and other vital characteristics. Evidently the preceding workload characterization. based on such vital characteristics, is not giving a complete description of this data. It is not sufficient to describe the distribution of the data points. It is also necessary to describe the sequence of those points.

R/S ANALYSIS

R/S analysis provides a means of accounting for the sequence of data points. Understanding this capability in a workload characterization context requires background in: fundamentals of fractal geometry, on which *R/S* analysis is built: origins of *R/S* analysis: and fitting *R/S* analysis to workload characterization.

FRACTAL FUNDAMENTALS

Basic understanding of fractals is built on two key concepts: fractal dimension and self similarity ([FEDEXX]). Feder adds that. "A neat and complete characterization of fractals is still lacking." (*ibid.*, p. 1 1), but these concepts suffice for a study of workload characterization.

FRACTAL DIMENSION

Difficulties in grasping fractal dimension usually arise from conflict with widely accepted aspects of Euclidean geometry. As indicated previously. fractal geometry accepts the existence of objects with dimensions that are not whole numbers. This is a major departure from Euclidean geometry. where objects must consist of all the points within the boundaries of the object. A one dimensional object. a line. consists of ALL the points between the two end points of the line. A two dimensional object. like a triangle. consist\ of ALL the points within the three boundary lines of the Figure. A triangle with some points missing is something less than two dimensional. but still more than one dimensional. It has a fractal dimension.

Noise patterns. like those in Figures 1 and 3. do not comprise all the points in a two dimensional space. They are. however. more than one dimensional. They have a fractal dimension. This makes it possible to use fractal geometry to assess the degree of fractal-ness directly, instead of trying to rationalize it away.

SELF SIMILARITY

Self similarity looks like a term that is self evident. In mathematical term\. it is not ([PEIT92], p. 161). There are excellent technical explanations of self similarity ([MAND77]. pp. 350.35 1 : (FEDEXX]. pp. 1 84- 1 89). In connection with workload characterization a simplified definition suffices. An object exhibits self similarity when parts of that object are statistically or geometrically identical to the entire object.

Statistical identity is the type of self similarity which pertains to *R/S* modeling. Simply stated, identity is established when the results of statistic\ applied to parts ot an object under study equal results of statistics applied to the entire object.

Selection of statistics to use in establishing statistical identity is a key issue. The preceding section on interpreting workload characterization questions the suitability of common statistical measures based on normal standard distributions and independent. identically distributed data in a workload context. The problem of accounting for the sequence of data points is not unique to workload characterization. Difficulties also arise in characterizing time series in areas like hydrology ([HURSS 1]) and economics ([MAND77], pp. 334-339). Effort\ to resolve these difficulties have resulted in a new analytical process and a new statistic that extends the usefulness of standard statistics: R/S analysis and the Hurst exponent.

ORIGINS OF ***R/S*** ANALYSIS

The history of *R/S* analysis is neatly summarized by Schroeder ([SCHR91]). It is built around the work ot Harold Edwin Hunt. Hurst's work characterizing the long term behavior of the flow of the Nile. among other thing\. has had considerable influence in the fields of chaos theory. dynamical systems. and fractal geometry ([MAND77], (SCHRYI], [PEIT92]). The focus of Hurst's work is on the cumulative effect of the variability in a system, rather than the moment to moment relationship of the state of a system to a central tendency. Another way of stating this is that

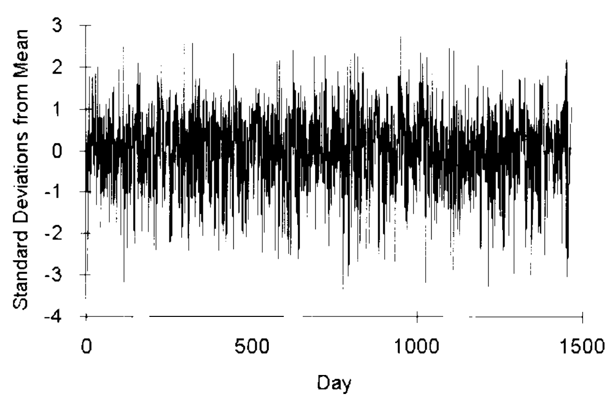

Figure 5. Randomly scrambled workload

Hurst studied the sum of the variations of system states from the mean. instead of considering each variation against others.

The impact of this approach can be demonstrated graphically using the real workload data. Figure 5 is a simple XY plot of the scrambled, z-scaled data over time. Figure 6 shows this scrambled data also, but as a cumulative series. The first data point is plotted at the mean. The second data point is plotted at the algebraic sum of its own z-scaled value plus the preceding plotted point. The third data point is. likewise, plotted at the sum of its own value plus the preceding (second) plotted point. Continuing the plot in this manner emphasizes the variability of the data. Figure 7 shows the real workload data in its original order as a z-scaled series.

There are two striking differences between Figure 6 and Figure 7. The most obvious difference is overall range. The vertical scale of Figure 7 is over six times that of Figure 6. The other major difference is smoothness. The line in Figure 7 has a less jagged appearance than that in Figure 6.

These differences share a common cause. The changes in Figure 7 tend to occur in the same direction more often than they do in Figure 6. If one value in Figure 7 is positive and increasing, the next value is likely to be positive and increasing, at least more often than is the case in Figure 6. Fewer changes of direction produce a smoother line. Successive changes in the same direction add up to greater range.

Cumulative variability in range is the basis for *R/S* analysis ([FEDE88], chap. 8). By comparing the cumulative variability of the subsets of a time series to the distribution of those subsets, ***R/S*** analysis provides a single statistic. the Hurst exponent. which can describe the differences between the scrambled and unscrambled workload data. The comparison process and meaning of the Hurst exponent can be explained using the workload data as an example and fitting the model to that data.

FITTING THE MODEL

R/S analysis is a simple but very repetitious process. The central activities are: accumulating the variations in a subset of data; finding the range ***(R)*** between the highest and lowest cumulative deviations calculated for the subset: finding the standard distribution (S) for the subset: and finding the rescaled range (*R/S*), which is the ratio of the range of the subset to the standard deviation ([PETE91], chap. 8). The central activities are carried out for subsets of the original data ranging in size from greater than one data point to less than or equal to half the size of the original set. *R/S* values of sets with the same size are averaged, thus yielding one ***R/S*** value for each subset size.

Applying this process to the workload data is straightforward. Taking an arbitrary subset size of five time ordered elements, one *R/S* value is generated for observations one through five; six through ten: eleven through fifteen: and so on through all the data points. Table 1 shows values which must be calculated for the first subset of five values.

These calculations are carried out for approximately 300 groups of five observations each in the workload data. The resulting ***R/S*** values are then averaged. producing a single ***R/S*** value of 2.109 for subgroups containing five observations.

The whole process is then repeated for larger and larger subgroups. A single ***R/S*** value is produced for each distinct subgroup size. At the final subgroup size. two, the *R/S* value will be the average of about 750 subset *R/S* values.

Comparing the change in *R/S* as the number of elements in a subset increases provides the Hurst exponent. Peters

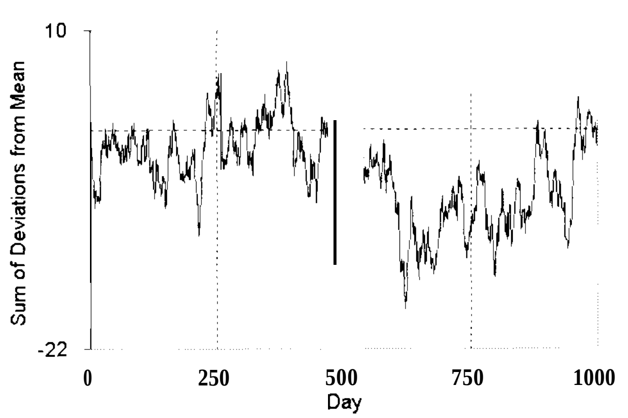

Figure 6. Cumulative plot of scrambled data

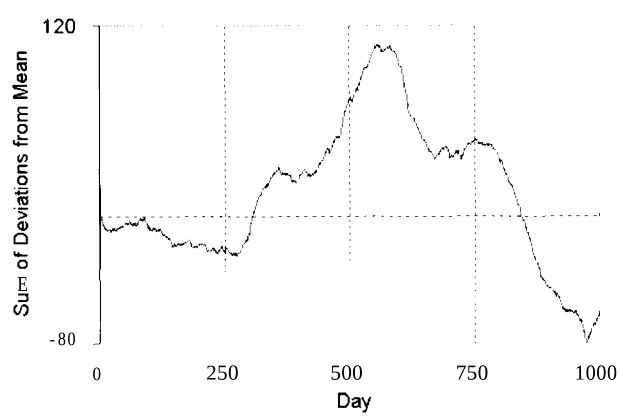

Figure 7. Cumulative plot of ordered data

([PETE91], p. 70) demonstrates that the Hurst exponent is related to the ratio of the log of R/S to the log of N. This means the Hurst exponent can be estimated as the slope of a regression line for the plot of $\log(R/S)$ versus log(N). The jagged lines in Figure 8 show a plot of this relationship. Fitting the elements into a linear equation gives us:

$$\log(R/S) = a + H \log(N) \quad (1)$$

where:

R/S = rescaled range

N = number of observations

a = a constant

H = slope of the plotted "line"

Appropriate regression lines and respective estimated slopes (Hurst exponents) also appear in Figure 8. The line labeled "$H = .85$" derives from ordered workload data. The line labeled "$H = .49$" comes from randomly scrambled data.

The reason for referring to the slopes in equation (1) as exponents can be seen by recasting equation (1) into:

$$R/S = a * N^H \quad (2)$$

The value of this recasting is that the resulting Hurst exponent provides a measure to characterize the differences in range between Figures 6 and 7. The Hurst exponent can also characterize the different in range within either Figure 6 or Figure 7 as the number of variations accumulated grows. This can be demonstrated by deriving the Hurst exponents for both sets of workload data and applying it to the display of the cumulative series.

The Hurst exponents in Figure 8 can be visually validated with the data in Figure 6. If this exponent characterizes the expected proportionate change in range as more variation is accumulated. then the Hurst exponent should determine the appropriate value to display reduced ranges of data at full scale. For example, to display half of the data points in Figure 6 at full scale. the expected vertical range is:

$$R_n = R_o / X^H \quad (3)$$

where:

R_n = expected vertical range

R_o = original vertical range

X = horizontal change factor

H = Hurst exponent

Number	Observed value	Deviation from the mean	Cumulative deviation
1	0.23597	.551956	0.551956
2	-0.2785	0.037486	0.589442
3	-0.02355	0.292436	0.881878
4	-1.32701	-1.01102	-0.12915
5	-0.18684	0.129146	**0.0**
Mean	-0.31599	**Maximum**	0.881878
		Minimum	-0.12915
		Range	1.011024
		Sample standard deviation	0.428
		R/S	2.362204
		Observations in subgroup	5

Table 1. **R/S** calculation for subgroups of 5 observations

Using scales from Figure 6 and the respective Hurst exponent from Figure 8 yields:

$$R_n = 32 / 2.49 = 22.8 \quad (3)$$

This indicates that a vertical range of 22.8 will handle 500 consecutive data points from Figure 6. Figure 9 shows this to be true.

Following the same process for the data in Figure 7. using a Hurst exponent of .85, gives a range of 111. which is pictured in Figure 10.

Applicability of this single statistic to all sizes of consecutive subsets within a time series is significant. It establishes the statistical identity that classifies the time series as self similar and fractal. In fact. [MAND77] and

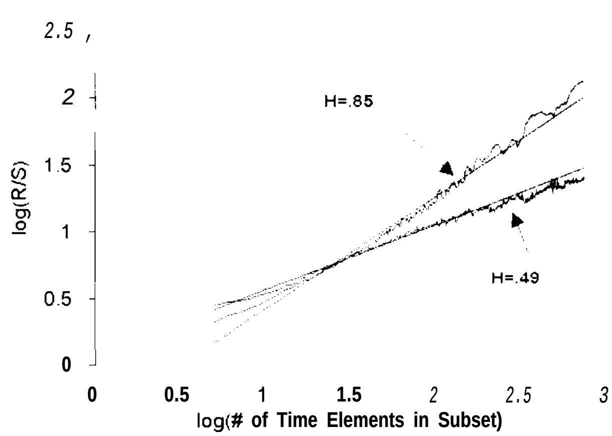

Figure 8. *R/S* **plot of ordered and random data**

Peters ([PETE91], [PETE94]) demonstrate clearly that, in time series like these, the Hurst exponent is the inverse of the fractal dimension of each series.

INTERPRETING R/S ANALYSIS

The hey to interpreting *R/S* analysis lies in the centrality of cumulative data to the process. *R/S* analysis. unlike most other statistical techniques, generates a statistic for a data point which is based on all the data preceding that point. Moving averages and exponential smoothing use some preceding data points in their processes, but *R/S* analysis uses all preceding data points. The unique ***R/S*** view of a time series enables it to characterize time oriented dependence in that series in ways which were not possible with more traditional techniques.

MEANING OF H

Mandelbrot ([MAND77], p. 386) summarizes the meaning of the Hurst exponent. ***H.*** First, the Hurst exponent, which is the resultant indicator from ***R/S*** analysis, can range between 0 and I. Midway in that range is a special value, .5. which indicates minimal timewise interdependence among the data points. It is convenient to think of this value as a traditionally random, or "SO/SO". chance that one data value will depend on another. As the Hurst exponent approaches I interdependence increases. a characteristic which is called "persistence." A persistent time series is one in which early values tend to have effects on later value\. causing them to vary in the same direction more often than is the case with truly random data. Conversely. data with a Hurst exponent below .5 is called "antipersistent." Long term dependence between data points causes them to vary in the opposite direction more often than randomness would dictate.

The two important factors here are long term dependence and likely direction of change. The first factor is particularly important. It quantifies the common sense idea that causes of change do not. necessarily, immediately precede the change.

IMPLICATIONS OF LONG TERM DEPENDENCE

Understanding the quantification of long term dependence requires some interpretation of Mandelbrot's summarization. It has been noted that a Hurst exponent of .5 indicates a lack of long term dependence in a time series. As long as the slope of the *R/S* regression line is not .5, a series has long term dependence. Change in the slope of the ***R/S*** regression line approaching .5, then. indicates that long term dependence is dissipating.

This applies to the actual workload data from Figure 3. Since the *R/S* plot for this data, in Figure 8, does not approach .5 within the time pictured. it can be assumed that long term effects in the workload data last longer than this time t'rame. On the other hand. the plot for the scrambled data labeled "H = .49" in Figure 8 shows that timewise interdependence was destroyed by the scrambling process.

This is interesting in the context of the original workload data. One example of an inference which can be made is that a transaction to enter a policy can have an effect on transactions entered years (hundreds of day\) later. Intuitively. this rings true. Once a policy is in force. there are certain business activities which must be carried out in support of that policy. Carrying this a little further. it seems obvious that a burst of new policies entered in a short time frame can produce a spike of support work that shows up years later.

SHORT TERM EFFECTS OF PERSISTENCE

Likely direction of change is the other main value provided by ***R/S*** analysis. This factor also derives from Mandelbrot's

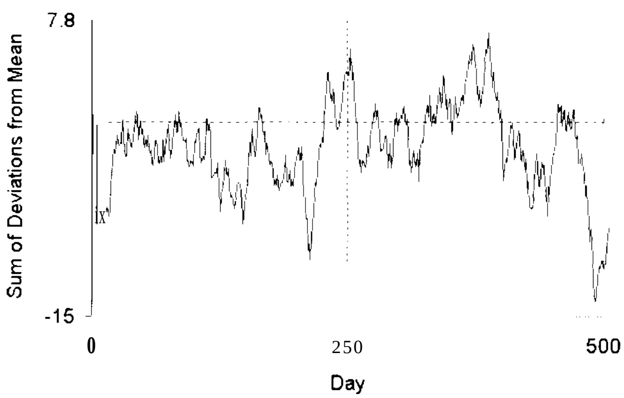

Figure 9. Scaled display of scrambled data

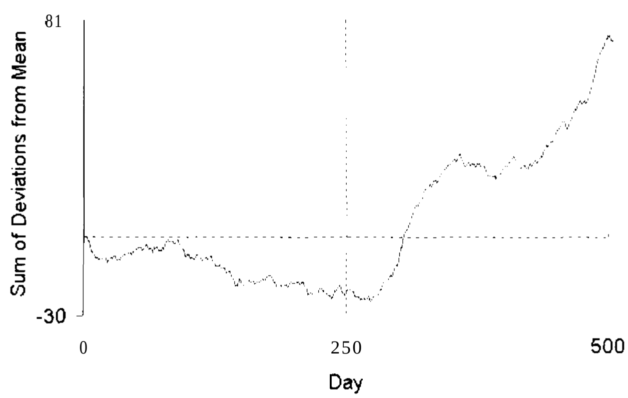

Figure **10.** Scaled display of ordered data observations on the meaning of the Hurst statistic. The relationship is very simple. The Hurst statistic provides a rough indicator of the percentage chance that a given data point will change in the same direction as its predecessor. Here again. it almost sounds like common sense when applied to the workload context. If more policies and data requests come in today than yesterday. it is reasonable to expect even more tomorrow. Once the load starts to decrease, it is reasonable to believe that there will be some "breathing room," and that the load will continue to decrease, for a while. The turning points are still pretty unpredictable, but activity between turning points is not as random as conventional methods indicate.

ALTERNATIVE FRACTAL APPROACHES

R/S analysis is investigated in this paper but other equally well conceived fractal models are available. The search for these techniques hinges on the status of fractal geometry as "... the geometry of chaos," ([PEIT92], p. 70). Searching for related terms like capacity planning. fractals, chaos. and dynamical systems, produces abundant evidence of the popularity of this area in business. Articles cover everything from general philosophical considerations in regard to the Russian economy ([MELL93]) to specific comparisons of chaos and other approaches for financial modeling ([NICH93]). Applications discussed range from stock market predictions to organizational analysis ([BAIL93], [WARN93]).

Interestingly, chaos techniques are almost unused in workload characterization. There are, in fact, only a small handful of reference\ that are relevant. Some furnish either too little detail ([ERWI99], [PICK88]) or too much ([HUBEXX]) to be useful in direct application to general workload characterization. The remaining articles tend to focus on individual components in a workload, like network queuing ([SMIT94]), disk access ([ERWI89]), or cache performance ([MCNU93]).

Shortage of prior work in high level workload characterization has not been a major problem. The wealth of material available in other disciplines provides alternatives. A striking similarity between these works is something Pickover ([PICK91], p. *vi*) calls an extension of "lateral thinking" which he maintains "...indicates not only action motivated by unexpected results, but also the deliberate shift of thinking in new directions to discover what can be learned." In other words, it is almost expected that the study of applying fractal to one field will be built on applications from other fields. Every major work on fractals/chaos theory/dynamical systems draws substance from many disciplines. Works on economics refer to works on psychology ([PETE91]), while works on psychology refer to works on computer science ([ABRA90]). These authors certainly think laterally.

Lateral thinking yields several alternative analytical techniques. Some involve concepts like bifurcation and strange attractors ([ROSS91]. [GOOD90], (MEDIC)'.?]). Explaining such concepts reduces the focus on clarifying the link between the fractal technique and the workload. By the same token, there is considerable appeal in approaches which emphasize more conventional time series characteristics. like range and deviation ([BROC91], [PETE91], (FEDEXX], [MAND77]).

The work by Brock, Hsieh, and LeBaron ([BROCY1]) is quite thorough. This thoroughness. unfortunately, diminishes the value of this work for a beginning approach. The included proofs and discussions of subjects like the generalized autoregressive conditional heteroskedasticity model are well outside the limits of this paper.

R/S analysis has been studied in this paper for several reasons. One key reason is the clarity with which]PETEY1] presented this approach as applied to an economic time series. Another reason is the abundance of reliable references on the technique ([PETEY1],]FEDEXX], [MAND77], [PEIT92]). Finally. the basic elements of the approach, range and standard deviation. are elementary concepts. Using these concepts allows explanation ot fractal analysis in more familiar terms.

RESULTS

Study of *R/S* analysis has produced positive results. It reveals fractal characteristic\ in the subject workload. It also provides a useful statistic that adds to currently available workload characterization. For example. the Hurst exponent is particularly interesting in the context of Kevin Smith's work (1994) on queuing of self-similar network traffic. Smith anchored his work on a fractal dimension of 1.66, which was drawn from a study on actual Ethernet traffic over several years (Leland. Taqqu. Willinger, and Wilson, 1993). This fractal dimension inverts to a Hurst dimension of about 0.6, falling in the same 0.5 to 1 range as the daily workload data analyzed in this paper. This commonality is important because Smith demonstrated that values in this range result in significantly higher queuing delays than would be indicated by more traditional analysis.

Another example of possible utility for the Hurst statistic is for simulation. As noted above, a true random number generator would create a simulated workload like that in Figure 5. It is visually obvious that such a workload contains a different combination of busy and slack periods from the actual workload in Figure 4. Several authors

([PETE91], [FINL93], [PEIT88]) present methods for generating an appropriate time series based on a principal called fractional Brownian motion. The methods of all of these authors depend on the Hurst exponent.

Study of *R/S* analysis has produced unexpected results too. The empirical nature of the approach has offered many challenges. Hurst formulated his analysis through experimentation on a time series which began in the year 622 A.D. ([SCHR91], p. 129). Experimentation still plays a big part in R/S analysis. Fitting *R/S* analysis to workload characterization entails considerable additional experimentation. Validation by analysis at different levels of granularity is one example of this. Another example is the coding of an analysis program to calculate the rescaled range. A third example is the continuous development going on in this field. Peters has just published a new work ([PETE94]) on fractal analysis which suggest refinements in determining what to model, and how to deal with factors which should be resolved before R/S analysis. Putting it briefly, the unexpected results relate to the amount of experimentation still needed, in spite of abundant references to R/S analysis.

ISSUES

The aim of this paper has been to demonstrate the value of fractal geometry to workload characterization. To achieve this aim is has been necessary to select a fractal model from among the many available in other disciplines: identify a suitable workload component on which to use that model: and evaluate the effectiveness of the model. While these objectives have been met. the results are not totally satisfactory.

R/S analysis is a widely accepted fractal process. Widely accepted. however. does not mean completely accepted. It is worth noting that [SPRO92] (p. 3X) use non-R/S analysis on some of the same data evaluated by Peters ([PETE91]) and assess the underlying structure to be simply random. Considerable skepticism remains about how clear the link is between chaotic. fractal models and econometric systems ([ROSS91], p. 120; [BROC91], p. 180).

Another source of dissatisfaction is application. It has been noted that major reasons for characterizing workload are understanding current resource utilization and predicting future needs. R/S analysis has been shown here to be useful for the first task: however. methods for prediction based on R/S results have not yet been developed. This limits the utility of R/S analysis as a characterization tool.

A third issue is interpretation. Since use of R/S is unprecedented in workload characterization. the focus here has been on the mechanics of generating the statistics with real workload data, not on explaining those statistics. Some possible interpretations have been suggested, but considerable work still needs to be done to attach concrete. quantifiable meanings to the results of R/S analysis.

CONCLUSION

In trying to answer one question, this paper has raised others. Three dominant questions are "What else might work?'. "What does this really mean?". and "What about prediction?"

There are, as noted previously. several approaches available. There are also other aspects of workload which might not work as well with R/S analysis as did transaction counts. More work is needed to assess the relative merits of adapting alternative fractal approaches to workload characterization.

Work is also needed in explaining the results. General observations about long term memory need to be extended to specify how long. and how strong that memory is relative to well understood process. There are similar needs for specific interpretations of short term effects as well.

The final main area of opportunity in R/S analysis is forecasting. If fractal geometry is only good for is describing what exists, workload characterization will be done with other tools. It is essential that workload characterization contribute to a picture of the future. drawn in inexpensive numbers. before critical decisions are needed on creating that future with expensive human resources. software, and hardware. Once forecasting is well integrated into the fractal toolkit. fractal geometry will almost certainly provide major enhancement to workload characterization.

REFERENCES

[ABRA90] F. D. Abraham. R. H. Abraham and C. D. Shaw. A *visual introduction to dynamical systems theory for psychology*. Ariel Press. Santa Cruz, CA (1990).

[ARMS82] R. M. Armstrong, *Capacity planning and performance management methodology*. IBM Corp., Gaithersburg, MD (1982).

[ARTI87] H. P. Artis, "Fundamentals of statistical workload forecasting." CMG *'87 Proc.*, 733 (1987).

[ARTI88] H. P. Artis and A. Sherkow, "Boosting performance with capacity planning," *Bus. Software Rev.* (May, 1988).

[BAIL93] L. Bailyn, "Patterned chaos in human resource management," *Sloan Management Review*, 34, 77-83 (1993).

[BORL92] Anon., *Borland Quattro Pro for DOS [computer program]*, Borland International. Inc., Scotts Valley, CA (1992).

[BORL93] Anon., *Borland C++ for OS/2 (computer program]*, Borland International, Inc., Scotts Valley, CA (1993).

[BROC91] W. A. Brock, D. A. Hsieh and B. LeBaron, *Nonlinear dynamics, chaos, and instability, statistical theory and economic evidence*. MIT Press, Cambridge (199 1).

[BRON80] L. Bronner, "Overview of the capacity planning process for production data processing," *IBM Systems J.*, **19, 20-21** (1980).

[ERWI88] H. R. Erwin, "Performance engineering techniques for complex dynamic systems," *CMG '88 Proc.*, l-4 (1988).

[ERWI89] H. R. Erwin, "Mixing and sensitive dependence on initial conditions in computer systems," *CMG Trans.*, 65, 3-6 (Summer, 1989).

[FEDE88] J. Feder. *Fractals*, Plenum Press, New York (1988).

[FINL93] M. Finlay and K. A. Blanton, *Real-world fractals*, M&T Books, New York (1993).

[FRIE88e] E. M. Friedman, "A workload characterization tutorial for on-line systems," *CMG '88 Proc.*, 920-931 (1988).

[FRIE88m] M. B. Friedman, "Nonlinear regression models and CPE applications". *CMG '88 Proc.*, 8-37 (1988).

[GAJ93] F. Gaj, "A business oriented capacity planning method: computer sizing through business volumes and indicators," *CMG Trans.*, *80*, 69-79 (Spring, 1993).

[GANT88] D. T. Gantz, "Tutorial on statistical methods for analysis and reporting for capacity planning," *CMG '88 Proc.*, *87* l-880 (1988).

[GOOD90] R. M. Goodwin. *Chaotic economic dynamics*. Oxford University Press, New York (1990).

[GORD83] P. A. Gordon. *So you want to estimate the value of availability*, IBM Corp.. Gaithersburg, MD (1983).

[HUBEXX] B. A. Huberman, (Ed.). *The ecology of computation*, North Holland, New York (1988).

[HUNT88] P. J. Hunter, "The capacity plan - table of contents," *CMG '88 Proc.*, *900-902* (1988).

[HURS51] H. E. Hurst, "Long-term storage of reservoirs," *Trans Amer. Soc. Civil Eng.*, **116,** 770-808 (1951).

[JAIN91] R. Jain, *The art of computer systems performance analysis: techniques for experimental design, measurement, simulation, and modeling*. John Wiley & Sons. Inc., New York (1991).

[KING86] G. King, "Workload characterization," *CMG '86 Proc.*, *789-80* l (1986).

[LELA93] W. E. Leland, M. S. Taqqu, W. Willinger and D. V. Wilson, "On the self-similar nature of Ethernet traffic," *Sigcomm '93*, Ithaca, NY (1993).

[MAND77] B. B. Mandelbrot, *The fractal geometry of nature (rev. ed.)*, W. H. Freeman and Company, New York (1977).

[MCNU93] B. McNutt, "A simple statistical model of cache reference locality and its application to cache planning, measurement, and control", *CMG '91 Trans.*, 13-21 (Winter, 1993).

[MEDI92] A. Medio, *Chaotic dynamics: theory and applications to economics*, Cambridge University Press, New York (1992).

[MELL93] G. Melloan, "On stop lights, chaos theory and Russia," *Wall Street J.*, A17 (May 3. 1993).

[NICH93] N. A. Nichols, "Efficient? Chaotic'? What's the new finance?" *Harvard Bus. Rev.*, **71, SO-**53 (1993).

[PEIT88] H. Peitgen and D. Saupe, (eds.), *The science of fractal images*, Springer, New York (1988).

[PEIT92] H. Peitgen, H. Jürgens and D. Saupe, *Fractals for the classroom: Part one: Introduction to fractals and chaos*, Springer, New York(1992).

[PETE91] E. E. Peters, *Chaos and order in the capital markets: a new view of cycles, prices, and market volatility*, John Wiley & Sons, Inc., New York (1991).

[PETE94] E. E. Peters. *Fractal market analysis: applying chaos theory to investment and economics*, John Wiley & Sons. Inc.. New York (1994).

[PICK88] C. A. Pickover. "Pattern formation and chaos in networks," *Comm. ACM*. **31**(1), 136- IS 1 (1988).

[PICK81] C. A. Pichover. *Computers and the imagination: Visual adventures beyond the edge*, St. Martin's Press. New York (199 1).

[ROSS91] J. B. Rosser. *From catastrophe to chaos: a general theory of economic discontinuities*, Kluwer Academic. Boston (1991).

[SAS85] Anon.. *SAS system for mainframes [computer program]*, SAS Institute. Inc.. Cary, NC (1985).

[SCHRY1] M. R. Schroeder. *Fractals, chaos, power laws. minutes from an infinite paradise*, W. H. Freeman and Company. New York (1YY1).

[SMIT94] K. F. Smith. "Effects of self-similar traffic on queuing systems." *CMG '94 Proc.*, 1236- 1247 (1994).

[SPRO92] J. C. Sprott and G. Rowlands. *Chaos data analyzer [Computer program manual]*, American Institute of Physics. New York (1992).

[WARN93] H. J. Warnecke, *The fractal company: a revolution in corporate culture*. Springer. New York (1993).