EUROPEAN COMMISSION
DIRECTORATE-GENERAL
Joint Research Centre

Institute for the Protection and Security of the Citizen

Time series analysis of European Grid Blackouts

E. Gutiérrez

European Laboratory for Structural Assessment (ELSA)
I-21020 Ispra (VA), Italy

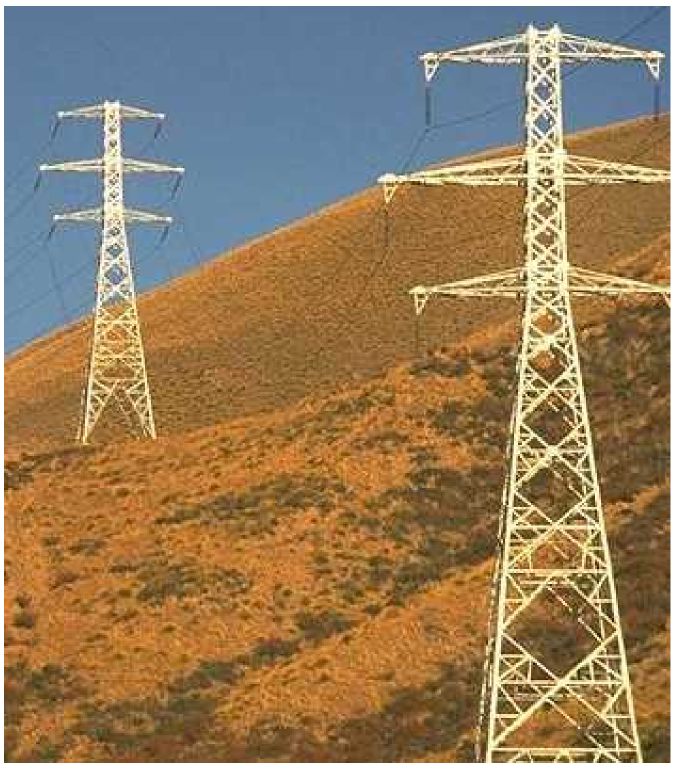

Institute for the Protection and Security of the Citizen

2005

EUR 21570 EN

European Commission
Directorate-General Joint Research Centre
Institute for the Protection and Security of the Citizen

Contact information
Address:ELSA TP 480
E-mail: eugenio.gutierrez@jrc.it
Tel.: ++-39-0332-785711
Fax: ++39-0332-789049

http://www.elsa.jrc.it
http://www.jrc.cec.eu.int

EUR 21570 EN
Luxembourg: Office for Official Publications of the European Communities

Printed in Italy

EUROPEAN COMMISSION
DIRECTORATE-GENERAL
Joint Research Centre

The mission of the Joint Research Centre is to provide customer-driven scientific and technical support for the conception, development, implementation and monitoring of European Union policies. As a service of the European Commission, the JRC functions as a reference centre of science and technology for the Union. Close to the policy-making process, it serves the common interest of the Member States, while being independent of special interests, whether private or national.

Time series analysis of European Grid Blackouts

Eugenio Gutiérrez
Institute for the Protection and Security of the Citizen
Joint Research Centre, European Commission.

Abstract

We analyse time series of electric power interruptions in the synchronous frequency network covered by UCTE (Union for the Co-ordination of Transmission of Electricity) over a period of 28 months. The scope of this preliminary study is to evaluate if the performance of this, the largest synchronously connected grid in Europe, shows characteristics of critically organised systems. In order to do so we have calculated the cumulative probability distribution and Hurst exponents of various blackout parameters. We have found indications that events in the time series do not conform to a normal statistical distribution, but rather, to a power law that implies that the distribution and magnitude of blackouts are correlated over substantial time-scales. The implications for the occurrence of large blackouts are discussed in terms of the availability of data sets over more extended periods of time.

1 Foreword

The issue of European electricity grid infrastructure security arises from our society's complete dependence on a reliable energy source to run the utilities necessary for everyday life. Demand and market forces have motivated the creation of a highly interconnected network that no single government or organization controls comprehensively. The stability and security of supply is measured by the capacity of the grid to react to unplanned perturbations of its normal working conditions. Such perturbations may result in local supply failures that may cascade into blackouts, or *outages*, that propagate over extensive portions of the system. Perturbations may originate from many sources, accidental or malicious, but their end effects and the methods of how best to deal with them will be similar.

The European electricity grid has, as a result of deregulation and competitive market forces, been running, on occasions, under conditions close to loss of equilibrium. This is because redundancy of supply and transmission capacity implies capital investments that cannot be offloaded onto a highly price-conscious market. Perturbations to the grid can rarely be predicted, however, the grid utilities have to deal with these in order to limit the propagation of cascading blackouts. To do so the transmission operators have developed sophisticated monitoring and control methods, and together they have willingly formed groups to develop regulations and procedures to mitigate the effects of disruptions.

EU directives such as those concerning deregulation of electricity-supply markets [7] on the one hand, or those addressing energy security issues [6] on the other, introduce changes to how the grid system is run. Such policy changes affect directly the physical performance of the system as a whole. For this reason it is believed that the European Commission, and associated bodies, will have to take a long term view with regard to the market and security policies it implements in this area. While it should ensure deregulation would benefit EU citizens, such policies should not result in mechanisms that might reduce further the stability of the grid.

A disruption of electricity supplies could have serious negative effects on the performance and security of the economy and citizens' everyday life. The security challenge is therefore to make the electricity grid infrastructure more secure without compromising the productivity and advantages inherent in today's complex interconnected networks. In order to do this both short-term and long-term technological development should result in the deployment of monitoring, *diagnostic* and control systems. These will affect some of the fundamental aspects of the organizations that run and supply the grid. In this paper we use the techniques currently being used to *diagnose* the qualitative nature of grid systems, and more specifically the statistical distribution of blackout events.

It is important to understand the generic underlying statistical nature of blackouts and outages that occur in European Electricity grid systems in order to assess the possible risk of large cascading blackouts. Our concern is not with the cause of the power interruptions (be they natural, accidental or malicious) but rather on how the system reacts to generic perturbations. The robustness of a grid system to absorb disruptions manifests itself in the tendency for it to damp out cascading blackouts. Clearly some outages are small and contained, others, such as the Italian blackout of 28th September 2003, affected practically the whole country for most of that day.

2 The UCTE system

The UCTE grid makes up the largest synchronously operating grid in Europe and includes both EU and non-EU member states (see Figure 1). Grid synchronicity allows the transfer of electricity from one state to another through a complex network of high-voltage lines that supply a population of 450 million people with an annual consumption of 2400TWh [10].

The main role of UCTE is that of establishing recommendations and technical rules for the synchronous operation of the UCTE grid. The aim of such rules is to render the UCTE system highly reliable addressing two main concepts:

- Security and reliability (short-term stability)
- Adequacy (long term supply and demand balance)

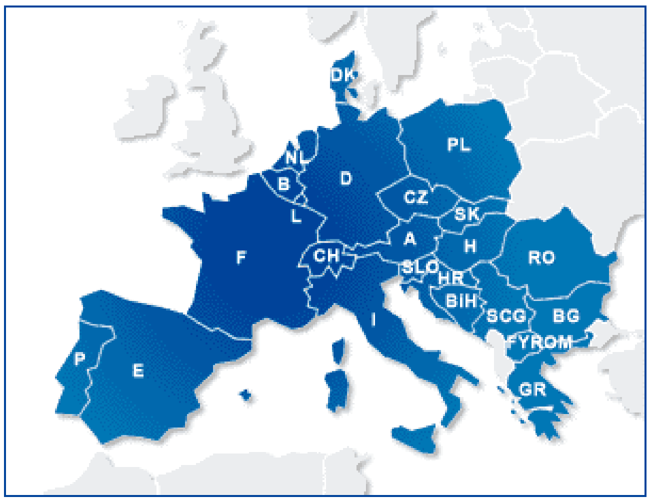

Figure 1 Member states of UCTE (Source UCTE)

UCTE has developed a number of rules and guidelines that constitute the basis for the smooth running of the grid. The operation of the system is therefore subject to security and reliability standards set within the framework of the UCTE cooperation agreement. One of the cornerstones of these guidelines is that the whole system must be capable of running normally even when a major failure occurs (e.g., the loss of a high-voltage line). This scenario is handled under the, so-called, N-1 rule which sets in place a series of measures aiming, in the first place, to ensure that stable operation is not disrupted and, secondly, to effect complementary measures to, at least, maintain the N-1 rule and avoid further cascading line failures.

In order to ensure that the N-1 rule is observed at all times, UCTE requires its members to monitor their systems closely so that, on the basis of grid performance data concerning any particular disruption, the operators can apply the remedial guidelines laid down. The operators monitor and control the performance of their grids in real-time with measurement devices placed at key locations. Grid performance data, such as voltages and currents) is analysed continuously in order to produce a state-space estimation of the system; this allows the operator to have a snapshot of their grid section ready, should an N-1 scenario occur. If a disruption occurs the operator can then take remedial actions in accordance with UCTE guidelines; quite often this involves exchange of information between neighbouring countries.

UCTE gathers grid performance data from its members and has recently started to provide regular reports on the performance of the grid. In particular, UCTE compiles lists of electricity power interruptions that have occurred in the synchronously connected grid system every month. These provide information concerning the time, location, magnitude and causes of incidents.

UCTE has provided the Institute for the Protection and Security of the Citizen of the Joint Research Centre, with a catalogue of power interruption data for analysis in order to qualify the statistical profile of the recorded events in its grid.

3 Electricity grids as complex systems

An analysis of the power interruptions on the UCTE grid may be used to examine under what conditions the grid has been operating; for example, if the system is running at the limit of its capacity, the probability of large blackouts[1] is more frequent than expected from standard statistical projections.

Recently, new methods of looking at grid blackout data [4] have been used to identify if such systems show indications of self-organised criticality: i.e., the tendency of certain, highly complex driving forces, to push the functioning of the system to conditions that may generate unstable behaviour. The underlying qualitative phenomena that drive coupled, spatially-distributed, networks can be found in a number of naturally occurring phenomena. For example, in analogy with snow or sand slopes, grids are subjected to ever-increasing demands that tend to load the system towards instability. Just as fresh snow piling up on a slope results in avalanches when the gradient reaches a critical inclination, the complex loading and dissipative mechanisms (that include market and regulatory forces, as well as technological developments), increase the demand on preferential grid lines, which increases the possibilities of large blackouts.

Cascading failures, just like avalanches, can be small or large, however, the distribution of these cascades conform to special statistical distributions which are the hallmark of complex self-organised critical systems (SOC). It is important for regulators and planners to understand the underlying nature and risk profile of SOC systems, so that they may put in place mitigation plans for extreme blackouts. In order to do so we must first establish if electricity grids behave like SOC systems.

A telltale sign of SOC systems is that the dynamics of certain parameters over time (or space) follow power laws. Expectations of large events do not decrease exponentially, but rather, do so in such a manner that the probability of extreme events may be orders of magnitude higher than expected. If a system as well regulated and monitored as the UCTE grid were to demonstrate signs of such critical self-organisation (as would seem to be the case in parts of the grid in the USA), then latching the present UCTE grid to the Baltic-Russian grid (generating a synchronously

[1] The term blackout used herein, refers to an electricity outage, or power interruption, of any magnitude ranging from 1MWhr to 177GWhr. In a scale-invariant system, all blackouts—irrespective of their size—are, in a phenomenological sense, identical. However for grid transmission operators a blackout usually refers to events affecting hundreds of thousands, or millions, of people. Herein we interchange the terms liberally, but we point out that, in an operational sense, the terms have significantly different meanings.

connected grid from Lisbon to Vladivostock) could, potentially, add considerable stresses that, if not properly considered, could make the whole system less stable than it is at present.

The presentation herein provides a preliminary analysis of the UCTE grid's power interruption data and is intended to gain some insight as to the statistical distribution of power interruptions and how events are correlated over time. The analysis does not purport to draft hard and fast conclusions on how the grid is operating but rather as a diagnostic method to identify the qualitative behaviour of the UCTE grid as a generic system.

4 Presentation of Data Set

The data set analysed in this paper was provided by UCTE. The events cover a period from January 2002 till March 2004 and include the Italian blackout of September 28th, 2003. For each event UCTE provides a comprehensive description of the main technical characteristics, including: the type of line power rating, country (location), elements (substations or transformers), as well as indications of the mechanisms that precipitated the blackout (natural or otherwise). In our analysis we have onlyconsidered the following parameters:

- time intervals between power interruptions.
- undelivered energy.
- power loss.
- duration of the event (restoration time).

Although the total number of events reported is 340, the data sets for the energy and power losses are incomplete. Some records contain no data for either energy or power loss, however, numerous examples contain one or the other. Thus, the data set for the undelivered energy is only 228 data points, that of the power loss was only 194, whereas the data set for the recovery times for each black-out is 333 points.

Some important details can be mentioned here. The longest wait between recorded power outages was 22 days, the shortest (other than synchronous events) is one minute. The smallest magnitude power interruption is 1MWh, the largest 177GWh (Italy September 2003). The smallest power loss is 1MW, the largest is estimated at just over 9GW (calculated from the energy loss and average duration given for the 177GWh Italian September 2003 blackout). The shortest recovery (duration of power outage) time is one minute, the longest 31 days. We shall discuss how we interpret recovery/duration time as it is not explicitly mentioned by UCTE; however, we presume that the duration time does not necessarily mean that a group of customers was not served for the exact time given; but, rather, the time during which a particular line was inoperative.

5 Analysis of Time series data

The time series data for the four parameters chosen for analysis are plotted as a function of time in days in Figure 2. The time interval between power interruptions, Δt, (Figure 2a), ranges from weeks to just a few minutes. The 2003 Italian blackout of September 28th is so predominant that it overshadows all other events in the energy and power-loss scales; hence we have also plotted the loss data in logarithmic scale in Figure 3.

A priori, from the correlation matrix, shown in Table 1, there does not seem to be any correlation between the magnitudes of the events—in terms of power or energy loss— and the time interval between events. Thus, based solely on these data, knowing something about the time intervals does not provide deterministic information about how big the next power interruption will be.

The power and energy losses, as expected, are highly correlated, but not equal to 1. One would expect that energy loss be trivially obtained by multiplying the power loss record by its corresponding restoration time. However, for those records where all three data parameters are available, the agreement is not one-to-one. In fact, as mentioned in §4, in some cases either energy or power loss data may be missing for certain records, for this reason we have chosen to analyse the energy and power loss data independently.

Regarding the restoration times, these appear to be uncorrelated to any of the other parameters. Thus, the magnitude of the outage and the time it takes to restore power, seem to be independent from each other. What this implies in terms of the operational characteristics of the UCTE grid is beyond the scope of this analysis, however, it is an interesting anecdotal point.

Table 1 Correlation values of blackout parameters.

	Δt	energy	power	$t_{restore}$
Δt	1.000	-0.045	-0.007	-0.091
energy		1.000	0.915	0.003
power			1.000	0.013
$t_{restore}$				1.000

6 Hurst exponent and system memory

Certain types of time series generated by natural processes may exhibit complex behaviour that may seem random but, in fact, are conditioned by deterministic phenomena. Such processes are said to possess long-term dependence (or memory).

Unlike normal Gaussian distributions, an event at time t_i has a direct effect on an event at time $t_{i+1\ldots n}$. The ***persistence*** of this memory can be measured with the Hurst exponent *H*. The Hurst exponent is a measure of the correlation of events within a time series and its value usually varies between 1 and 0. Thus, for deterministic signals $H = 1$, for purely noisy signals, as is the case for a Gaussian distribution, $H = 0.5$.

The Hurst exponent is estimated by sampling a complete data set over multiple time-window frames, and then calculating the, so called, average re-scaled range (R/S) for each window. The slope of the plot (in log-log scale) of the R/S values versus the window size gives the Hurst exponent. A full description of how the Hurst exponent was evaluated for the UCTE data set is given in the Appendix.

When $1 > H > 0.5$ the time series is said to have long-term memory, and the data are positively correlated in some manner. Thus, some randomly generated events will have an effect on future events that is forgotten with time. Events in the future are positively correlated to those in the past; hence increments (respectively, decrements) will be followed by increases (respectively, reductions) of the measured variable at some time in the future. The higher the *H* value is in this range, the more persistence it is said to have.

When $0.5 > H > 0$ the time series is said to be ***anti-persistent*** —or anti-correlated— in the long range. Thus, for anti-persistent time series, increments will be followed in the future by decrements (or the converse: decrements will be followed by increments). In natural phenomena persistence is more common, however, in some complex networks (such as financial markets) anti-persistence is exhibited for some time scales and is associated with ***volatility***. In the economics literature, anti-persistent processes are referred to as ***mean-reverting***; i.e., a process whose mean is zero over a sufficiently long time scale. Thus, financial markets and related systems [11], such as the daily bidding price in electricity markets, are often punctuated with periods of high and low volatility which is a sign of anti-persistence.

7 Analysis of Hurst Exponents

We examine the Hurst exponent of each parameter in turn.

7.1 *Time between blackouts:*

The trend in the intervals of the blackout data has an average slope of approximately 0.75 over the whole range of time lags. This implies that the intervals between reported power interruptions (irrespective of their magnitude in terms of energy or power loss) are highly persistent. Thus if a short time interval occurred between two blackouts it will have a persistent effect over the whole time span. The reverse could also be true; i.e., the effects of large intervals between successive power interruptions will persist over the entire range. The conclusion is that the time intervals between successive power interruptions are persistent and correlated.

Had the slope been equal to 1 (*H*=1), then all events would follow deterministically from the very first event. Conversely, had *H*=0.5 all the events would be independent from each other. In our case, given that $H \approx 0.75$, we must assume that some events are random, and many are the deterministic result of prior ones. Thus, some blackouts are randomly generated but these, in turn, generate others, which is a paradigm for cascading phenomena. The extent of the cascade cannot be observed from this graph, we can only conclude that many events have a deterministic relation to prior events over the time-scales studied.

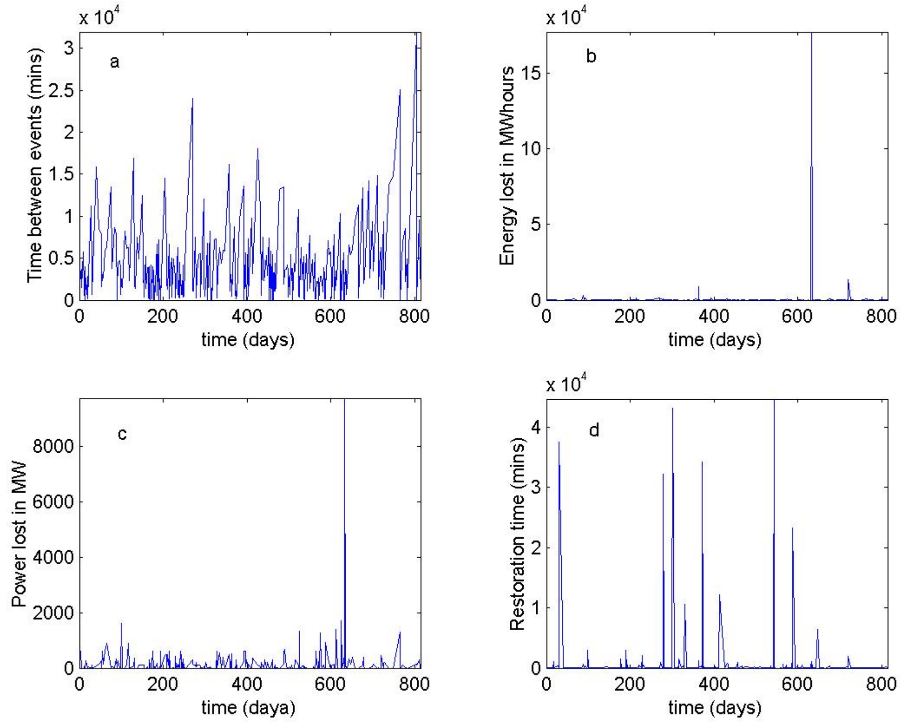

Figure 2 Time series of power interruption data from UCTE grid.

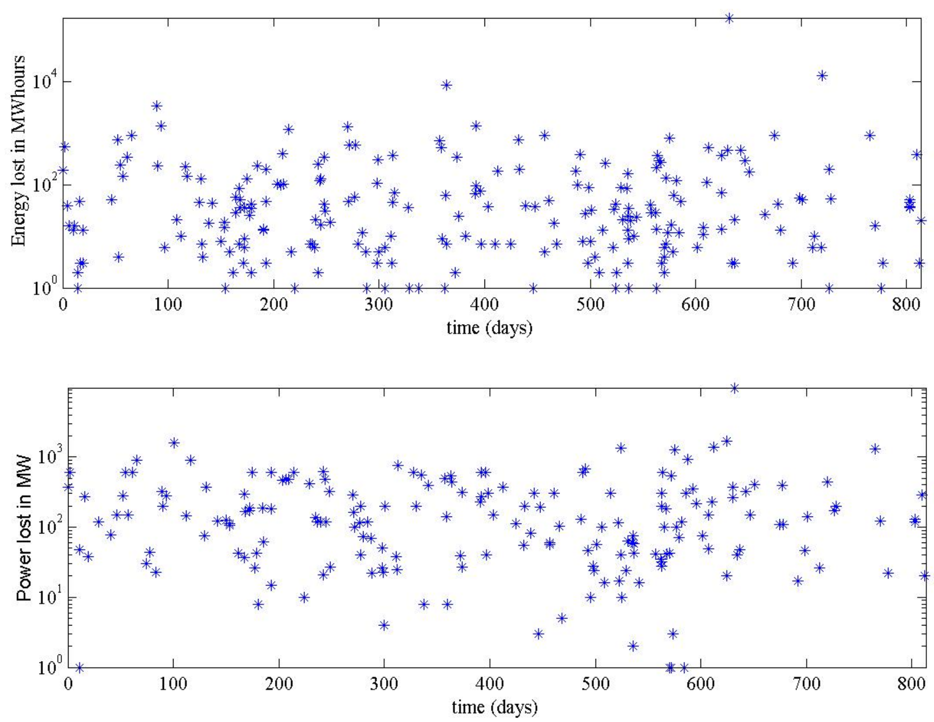

Figure 3 Detail in logarithmic scale of Energy and Power loss.

7.2 Magnitude of blackouts: Energy and Power losses

Whereas the R/S data for the time interval between blackouts could be approximated well by a simple straight-line fit, the situation is markedly different for the energy and power loss data. In the range of time windows corresponding to 1 to 5 months, the data for the energy loss can be fitted with a slope corresponding to $H \cong 0.64$, whereas for longer time windows (5 months to over two years) the slope is $H \cong 0.35$.

From the power loss data we observe the same type of behaviour corresponding to similar time windows; thus, in the range 1 to 5 months we have $H \cong 0.67$, which is indicative of high persistence, whereas for longer ranges we have $H \cong 0.44$, and hence, anti-persistent. We conclude that, in the short-to-medium time window ranges, the size of a blackout will have a noticeable persistence in the future up to periods of half a year or so: i.e., small (respectively, large) blackouts will be followed by smaller (respectively, larger) blackouts in the short to medium term.

In analogy to the financial markets, these data can be understood in terms of trends. Thus, a downward trend is peppered with instances where the market value of a range of stocks increases—maybe even substantially—within a time window; however, the overall stock value has effectively decreased. Fluctuations within a downward or upward trend are purely random elements superimposed on a deterministic signal. Of course we cannot predict what will happen in the future, as our interpretation can only be done with hindsight; but, what we can say, is that within periods of six months, the magnitude of the power interruptions in the UCTE area were not randomly distributed.

So, what size of blackouts can we expect within a three-month period if we have data for the last three months? We could of course conduct a statistical risk analysis on the basis of the acquired data; but, in order to do so, we need to establish what kind of statistical processes are driving the dynamics. Clearly, the present data set does not conform to a purely random process due to the fact that $H \neq 0.5$ so we must have some idea of the statistical distribution of events, in particular, extreme events. We shall deal with this in the following chapters when we analyse the cumulative distribution functions of the data set.

However, first we turn back to fact that the Hurst exponent for the longer time windows is less than 0.5. As mentioned above, Hurst exponents less than 0.5 are indicative of anti-persistent behaviour, which is associated with volatility —a mean-reverting process whose average over the period in question tends to zero. In this range, the system dynamics varies wildly and strong anti-correlations predominate: i.e., large blackouts will be followed by small ones and vice versa. The primary mechanisms driving the system over long time windows may be very different to those that generate persistent behaviour in the short-to-medium term. It is presumed that some form of strong feedback in the system drives it in a chaotic manner.

Anti persistence has been identified in financial markets and, recently, such phenomena have been studied in the time series of the *Nordpool* electricity spot market price [11]. The fact that it also appears in the time series for blackout size is noteworthy; but, it would be premature (based on this small sample size) to insinuate that the type of market and regulatory mechanisms that drive volatility in the *Nordpool* market may also drive the size of blackouts in the UCTE area. For the moment we suggest that further study of such correlations (for example, studying the volatility of market prices in the UCTE area for the same period examined in this study) may be worth pursuing.

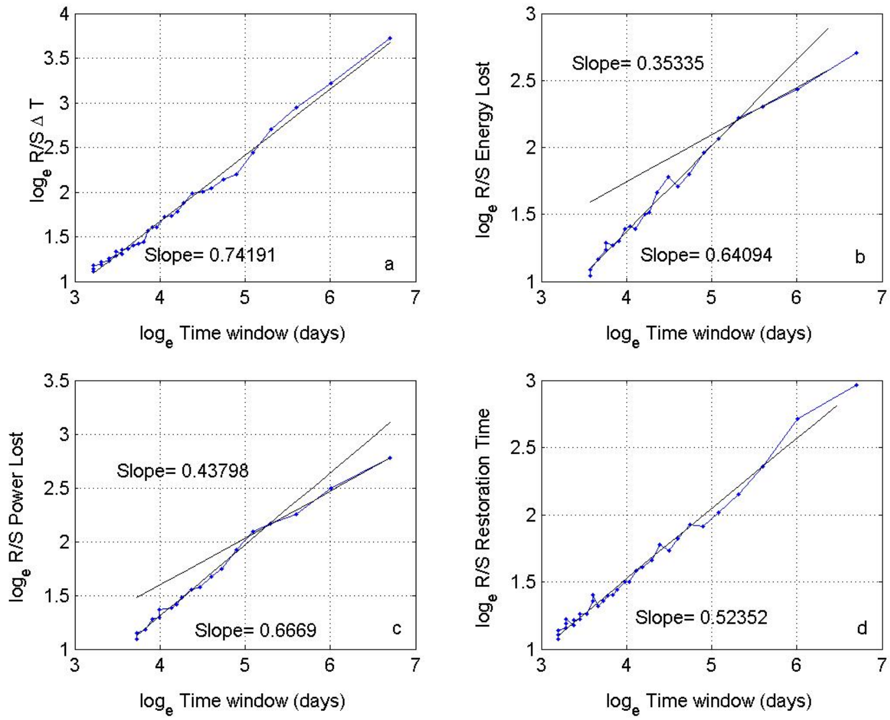

Figure 4 Hurst exponents obtained from rescaled time series data.

7.3 Restoration time

In contrast to the energy and power loss data (but as was the case for the time between intervals), the restoration time data can be well fitted by a single straight-line plot. However, whereas the R/S slope for the time intervals clearly showed signs of persistence, here we can assume that the restoration times approximates well to a pure random process because $H \cong 0.52$, is closer to what we expect for a random distribution. In fact, comparing the size of the power interruptions to their restoration times, we notice no clear correlations: some large events being dealt with relatively quickly and other, smaller ones, taking much longer. We can assume that, to some extent, the restoration times are driven by different factors to those that cause the outages in the first place. However, there can be no restoration time without a blackout, so some, *de facto*, causal relation exists between them. Restoration procedures, it is presumed, vary from country to country; restoration practices do not form a grid system (but are connected to it indirectly) and hence cannot exhibit the same type of persistence across time scales. Much depends on what is meant by *restoration time*; but, at present, this subject is outside the scope of this analysis.

8 Cumulative frequency and probability density distributions

The evolution of many ergodic signals can be described in terms of their spectral distribution by a power law inversely proportional to the frequency. So-called power laws occur in natural phenomena [2][8], and are said to be scale invariant.

When examined in the frequency domain such processes are said to have power at low frequencies. We shall discuss the implications of this below. Generically, a signal whose power spectrum decays as a function of $1/f^{a}$, does not have a normal, Gaussian, statistical distribution. Thus, for random white noise $a=0$, for Brownian motion $a=2$, whereas for generic *1/f* noise $a \cong 1$.

When a dynamical system exhibits a parametric power law distribution, it is a suggestion that long-range correlations between widely varying time (or spatial) scales of the system are occurring [1]. Such critical systems also exhibit scale invariance and fractal structure. However, critical systems are inherently unstable: i.e., arbitrarily small perturbations can change the value of a critical parameter that generates a qualitative change in its performance. A self-organised critical system is one that evolves in such manner that leads it to operate at conditions close to criticality, hence the term *self-organized criticality* [1].

In SOC systems, the tail of the probability density distribution when plotted in log-log scale, decays in an algebraic manner, so that the probability of occurrence of large disruptions decreases as a power function of the event size, whereas for processes characterised by normal random distributions, the probability of large events decays in an exponential manner. Stated simply the probability of large catastrophic disruptions in SOC systems are more than would be expected by extrapolating a normally distributed probability function. In terms of risk evaluation, the negative effects may be compounded by the fact that the actual cost of blackouts does may not grow linearly with the blackout size [4].

In this paper we shall not analyse the tail of the probability distribution function (PDF) as its evaluation is susceptible to noise and requires a well populated data set. Instead, we shall analyse the tail of the complementary cumulative frequency distribution (CFD), which is simply the integral of the PDF. In this case, the slope of the PDF in log-log scale is given by

$$\alpha_{PDF} = \alpha_{CFD} - 1.$$

The complementary cumulative frequency distribution is obtained trivially by counting the number of samples whose magnitude is less than, or equal to, a given value of parameter being examined. A follow-up paper will compare CFD and PDF slopes calculated on the same data set presented herein.

9 Analysis of Cumulative frequency

The cumulative frequency data for the four parameters studied is shown in Figure 5. The distribution of the time intervals between power interruptions (Figure 5a) is quite flat up to a range of 2^{10} minutes (i.e. less than 24 hours). For long waiting times, the slope in the last portion of the cumulative frequency curve is quite steep ($\alpha_{CFD} \cong -2.15$, corresponding to a slope in the probability distribution function $\alpha_{PDF} \cong -3.15$); thus, time intervals between outages longer than a few days are a rarity. Conversely, the expectation of hourly outages (irrespective of their magnitude) somewhere in the UCTE grid is quite high.

The frequency distribution of the magnitudes of the blackouts is significantly different from their time intervals. From Figure 5b, we see that the rate of decay is much shallower, having an average slope of $\alpha_{CFD} \cong -0.85$ over a wide energy-loss scale. In fact, this slope would have been lower had we included the country-wide September 2003 blackout in Italy. What this graph implies is that the probability of large blackouts would seem to follow a power law with probability distribution function slope $\alpha_{PDF} \cong -1.85$. Hence, we can expect that a normal Poisson process does not govern the probability of large blackouts, and that very large blackouts would be more frequent than expected. In fact, the large Italian blackout of 2003 is even in excess of the value expected from the power law fit; i.e., had this blackout been included, the slope would have been even shallower.

The plot of the CFD for the power loss, Figure 5c, is not as pronounced as that of the energy loss; but it too, is also indicative that a power law governs the probability distribution of large power losses.

Given that the ratio of power and energy losses provided by UCTE is not identical to the restoration time (i.e. energy loss is not trivially obtained by multiplying power loss by the restoration time), we cannot expect their respective cumulative frequency distributions to be identical.

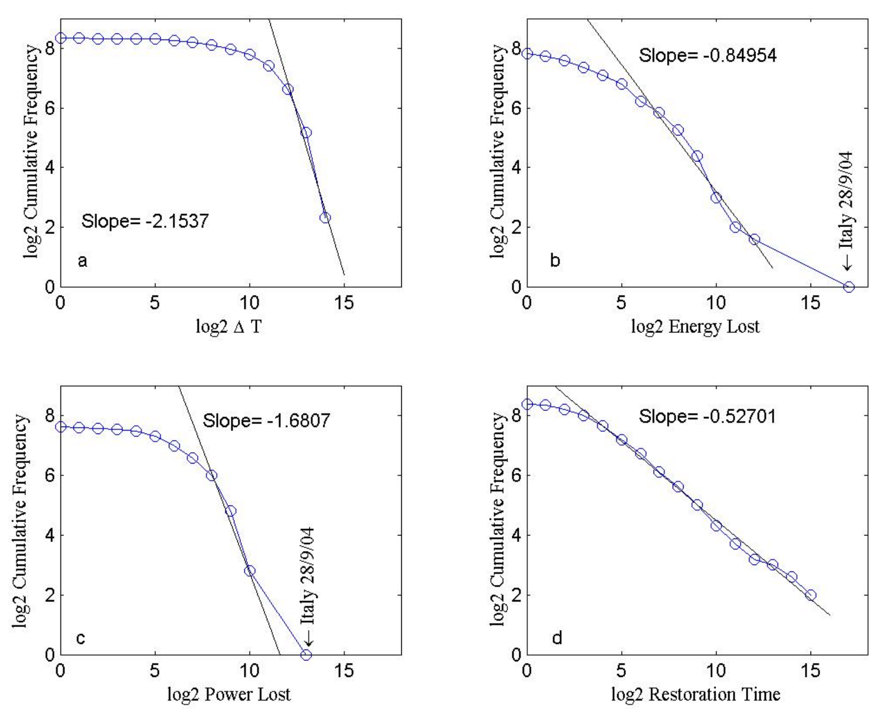

Figure 5 Cumulative frequency plots.

Concerning the restoration time itself, we find that its frequency distribution, shown in Figure 5d, in analogy with the energy-loss, also has a shallow slope. One could assume that energy and restoration times are more closely correlated; however, such an assumption should consider what is meant by *restoration time*; i.e., whether it is the amount of time that energy was not effectively delivered to a customer or, mechanically speaking, the time it took to repair a particular line whilst still being able to deliver power to the customers in the area affected. Irrespective of these unknown factors, it can be said that the slope of the restoration time is also governed by a power law (slope –0,52 and –1.52 for the cumulative frequency and probability distributions respectively). This slope implies that very long restoration times in the system may occur, and that if these are linked to wide sectors of the grid, substantial blackouts are highly probable.

In view of the limited data set, one should approach conclusions carefully. For example the Italian 2003 blackout lasted approximately 18 hours (although substantial parts of the grid recovered after only a four hours), however, the size of the grid affected was so large that the resulting energy loss was the highest in the UCTE data set analysed herein. As a general rule it can be said that, the expectation of outage times are higher than would be for a normal random distribution, hence expected energy losses are, statistically, in proportion.

Conclusions

An analysis of power interruption data from the UCTE grid system reveals that the time intervals between blackouts is not a purely random process; but, rather, events in the future are conditioned deterministically to some precursor (possibly random) event in the past (this is borne by the fact that the Hurst exponent, $H \cong 0.74$, is indicative of persistent behaviour). As a consequence, we expect that short intervals between electricity outages will be followed by even shorter ones in the near future and we postulate that, when many power interruption intervals coalesce, they generate cascading blackouts. The size of the cascade will then determine the magnitude of the blackout. Such qualitative behaviour is symptomatic of self-organised systems and analogous to sand-pile and snow avalanches.

Energy and power losses over short to medium periods (from a few days up to six months), also seem to be persistent $H \cong 0.64$. For these time-scales small or large blackouts will be followed by, respectively, smaller or larger blackouts in the future. However, for longer time scales (from six months up to two years) blackout magnitudes show signs of anti-persistence $H \cong 0.35$, whereby small —respectively, large— blackouts will be followed by larger —respectively, smaller— blackouts. This is indicative of *volatility*, which may be associated with analogous behaviour in some European electricity price markets.

Whereas time intervals, energy, and power losses of blackouts show signs of persistence (or anti-persistence as the case may be), the data for the restoration time between events seems to have no time correlations: i.e., duration of outages are independent at all time scales on the UCTE grid. In spite of this, the corresponding cumulative frequency distribution seems to follow a power law, which implies that the expectation of large restoration times is higher than would be expected for a normally distributed random process. The same argument applies to the cumulative frequency distributions of energy (and to a lesser extent) the power loss. For these we find that the expectation of large blackouts can be fitted by a shallow-slope power law; thus, the expectation of large, country-wide, outages such as the Italian blackout of September 2003, is even under predicted by the power law fit obtained by excluding that particular event from the calculations.

Although these data indicate that the UCTE grid system (in analogy to other electricity grid and complex networks) shows signs of self-organised-critical behaviour, the results presented in this paper cannot be regarded as definitive due to the dearth of data used in the analysis.

The conclusion is not that the UCTE grid is on the verge of catastrophic failure, but rather, that to increase its security further still, the use of non-linear analysis techniques could be applied in order to develop diagnostic systems more suitable to such complex distributed systems. Such diagnostic techniques, used in conjunction with innovative control hardware implementations, such as using Phase Monitoring Units deployed at key positions on the grid, will most probably develop over the forthcoming years as a means of improving the security and reliability of the EU electricity network. However, due to the interconnectedness of the grid, such systems will only be effective when conducted in collaboration with all the connected transmission system operators.

Acknowledgements

The author would like to thank F. Cahyna, O. Feix and J. Fijalkowski, from UCTE, for their collaboration and providing the documentation and data sets that have been used in the analysis presented herein. Thanks go also to Y. Drossinos and J.M. Zaldivar, from the Institute for the Environment and Sustainability of the JRC, for the helpful discussions and comments on the analysis.

References

[1] P. Bak, C. Tang, K. Wisenfeld**, *Self-organised criticality*,** Phys. Rev. A, 1988**, 38 ,** 364-374.

[2]G. Boffeta et al.**, Power laws in solar flares: self-organised criticality or turbulance?**, Phys. Rev. Letters**,** vol**. 83, 22**, 1999, 4662-4665.

[3] M. J. Cannon et al.**, *Evaluating scaled windowed variance methods for estimating the Hurst coefficient of time series,*** Physica A: Statistical and theoretical physics**,** vol.**241**, 1997, 606-626.

[4] B. A. Carreras. D. E. Newman, I. Dobson, I. B. Poole**, *Evidence for self-organized criticality in a time-series of electric power system blackouts,*** to appear in IEEE Transactions on Systems and Circuits Part 1.

[5] P. Devynck, G. Wang, G. Antar, G. Bonhomme, ***The Husrt esponent and long-time-correlation,*** 27th EPS Conf. On Cotr. Fusion Plasma Phys. Budapest, ECA Vol. **24B** ,2000, 632-635**.**

[6] Energy Infrastructure and security of supply, COM(2003)743 final, European Commission, Brussels, 2003.

[7] Concerning common rules for the internal market in electricity, Directive 2003/54/EC, European Commission, Brussels, 2003.

[8] P. Hähner and Y. Drossinos, ***Scale invariance and intermittency in a creep-slip model of earthquakes***, Phys. Rev. E, vol **59, 6,** 1999, 6231-6233.

[9] H. E. Hurst**, *Long-term storage capacity of reservoirs*,** Trans. Am. Soc.Civil Eng., vol. **116**, 770, 1951.

[10] Union for the Co-ordination of Transmission of Electricity, http://www.ucte.org/

[11] I. Simonsen, ***Measuring anti-correlations in the Nordic electricity spot market by wavelets,*** Physica A: Statistical Mechanics and its applications, vol. **322**, 2003, 597-606.

Appendix

Evaluation of the Hurst Exponent

The derivation for the Hurst exponent [9] presented herein is as follows:

Given a discrete time series $x(t)$ ranging over a period τ whose mean is given by $\langle x \rangle_\tau$,consider a point of the data series at time t_n, and define its deviation from the mean as $x(t_n) - \langle x \rangle_\tau$. We perform this calculation for each point of the data series and sum over the complete range τ thus:

$$X(t,\tau) = \sum_{t_n=1}^{\tau} x(t_n) - \langle x \rangle_\tau \ .$$

So $X(t,\tau)$ is the cumulative sum of the deviation of the signal from the mean. The range of $X(t,\tau)$ is defined as $R(\tau) = Max\left(X(t,\tau)\right) - Min\left(X(t,\tau)\right)$.

The rescaled range is obtained by dividing the range by the standard deviation of the signal, hence:

$$R/S = \frac{R(\tau)}{S(\tau)},$$

where $S(\tau)$ is the standard deviation of the range of $X(t,\tau)$ over the period τ which is simply given by

$$S(\tau) = \sqrt{\tfrac{1}{\tau} \cdot \sum_{t_n=1}^{\tau} \left\{ x(t_n) - \langle x \rangle_\tau \right\}^2} \ .$$

In order to evaluate the Hurst coefficient we monitor the evolution of the R/S data as a function of the sample window size. Thus, the first point of the R/S data corresponds to window size w_τ covering the whole period τ. We then proceed to analyse the data in segments of ever-decreasing size $w_{\tau/n}$. For example, a non overlapping bisection of the original time window would produce two R/S values corresponding to the two $w_{\tau/2}$ windows covering the first and second halves of the whole range. We then take the average of these to obtain a single R/S value corresponding to a window size $\tau/2$. We proceed with more bisections until we deem that the number of points of the window size is too short to provide significant data. There is no hard and fast rule that defines the size of the smallest window, and some authors have analysed this aspect as a function of the length of the time series [3],[5]. In general we should consider the sampling frequency and the inherent dynamics of the event (if we know something about the time scales of the underlying driving process). Likewise, there are no specific constraints as to the rate of decrease of the window size or their placement along the time period (i.e. shifting window segments may or may not overlap). For our analysis we have chosen sequential windows whose size grows as a linear function of the signal range.

The Hurst exponent is the slope of the $\log_e$-$\log_e$ plot of the R/S versus $w_{\tau/i}$.