Network Traffic Measurements in a Switched Ethernet Environment

Tony Field, Uli Harder & Peter Harrison *

June 16, 2002

Abstract

We present new measurements of network traffic in a switched Ethernet system and file size/webserver request size distributions. We find that the distribution of packet rate changes can be well described by a (truncated) Cauchy distribution. Similarly, for the main contributors to the network traffic, file and web document requests, we can show that the distribution is similarly described by a (truncated) Cauchy distribution.

1 Introduction

In this paper we present some of the results of an analysis of computer network traffic that we have measured in the Department of Computing of Imperial College. The measurements were performed in order to parametrise models of the network traffic realistically.We are going to develop later.

There have been many studies into network traffic behaviour since the seminal papers by Leland et al. and Eramili et al. [2, 3]. The main feature they found was that network traffic exhibited self-similar features which could not be explained with simple Poisson processes. Since then many explanation of the origin of the long range dependence and models for it have been put forward. They reach from ON/OFF models of heavy tail distributions [2, 3], through the file size distribution in file systems and web servers [4, 5] to user behaviour, back-off algorithms in the Ethernet [7], higher level network protocols, buffers in routers and the TCP congestion avoidance algorithms [8, 9, 10, 11].

In a previous publication [12] we have shown that the network we measured shows the usual features of self-similarity in the power spectra. In this paper we try to investigate the relationship between the distribution of file sizes and requests to a webserver with the changes in the packet rate of the network traffic we measured.

File sizes and request sizes are well known to exhibit power laws [4, 5]. And there have also been attempts to explain why we see these power laws [13]. We find that in our measurements the distribution are well approximated by the Cauchy distribution that [13] predict. We also find that the change of packet rate distribution is well approximated by a truncated Lévy distribution. This type

*Department of Computing Imperial College of Science, Technology and Medicine Huxley Building 180 Queen's Gate London SW7 2BZ England {ajf|uh|pgh}@doc.ic.ac.uk

of measurement is commonplace in the analysis of stock prices, for example [14], but not, to date, for network traffic.

The measurements we have performed are cheap and easy to reproduce at other sites. In particular, in our discussions we promote the use of the data in the virtual file `/proc/net/dev` file when analysing over very long time periods. This is both economical on resources and can be done without system privileges.

The rest of the paper is organised as follows: We describe the way data has been captured and the architecture of the network in question in section 2. In the next section 3 we describe how the data was analysed and present our findings. The last section 4 is indicating future research directions, with special emphasis on simulations and the use of an adapted form of the diffusion approximation.

2 Network Architecture and Monitoring

The monitored system is a departmental switched Ethernet. We focus attention on three components of this network for the purposes of this paper, as illustrated in Figure 1: The router, which connects the network to the outside world, an arbitrarily chosen CPU server (named MOA), and the departmental web server which services both internal and external web page requests.

All the relevant data that has been gathered in this project has been made available to the public [15]. In this paper we will concentrate on the data collected from the `/proc/net/dev` files, the webserver logfiles and file size measurements. In this paper we will not discuss the `tcpdump` [16] results, which have investigated in details in [12].

2.1 File and Request Sizes

We have chosen a 2 hour period on 22 March 2002 between 12pm and 2pm for our detailed investigations. At the application level we used the logfile created by the Apache server. The logfile keeps a record for each request made to the web server and logs the number of bytes shipped to the client. Here we concentrate on the log file sources as we want to determine the request size distribution. We filtered out internal request as they do not find their way into the router monitoring files.

In addition we have verified on a number of Solaris and LINUX machines that the file size distribution of files on the local disk is indeed distributed according to a power law (using an adapted script from [4]).

2.2 Packet Rates

For measurements at the network level we made use of the file in the `/proc` file system `/proc/net/dev` where the Linux operating system keeps track of the number of packets and bytes sent using counters. The counters get reset at machine reboot, the values are modulo 2^{32}. We have used a `PERL` program to query this file at regular intervals of 1 second and record the value of the counters and a time stamp. Occasionally this process fails to record the counter within 1 second. In the processing phase we then linearly interpolate the missing values and get a time series of the counter values for every second. The data

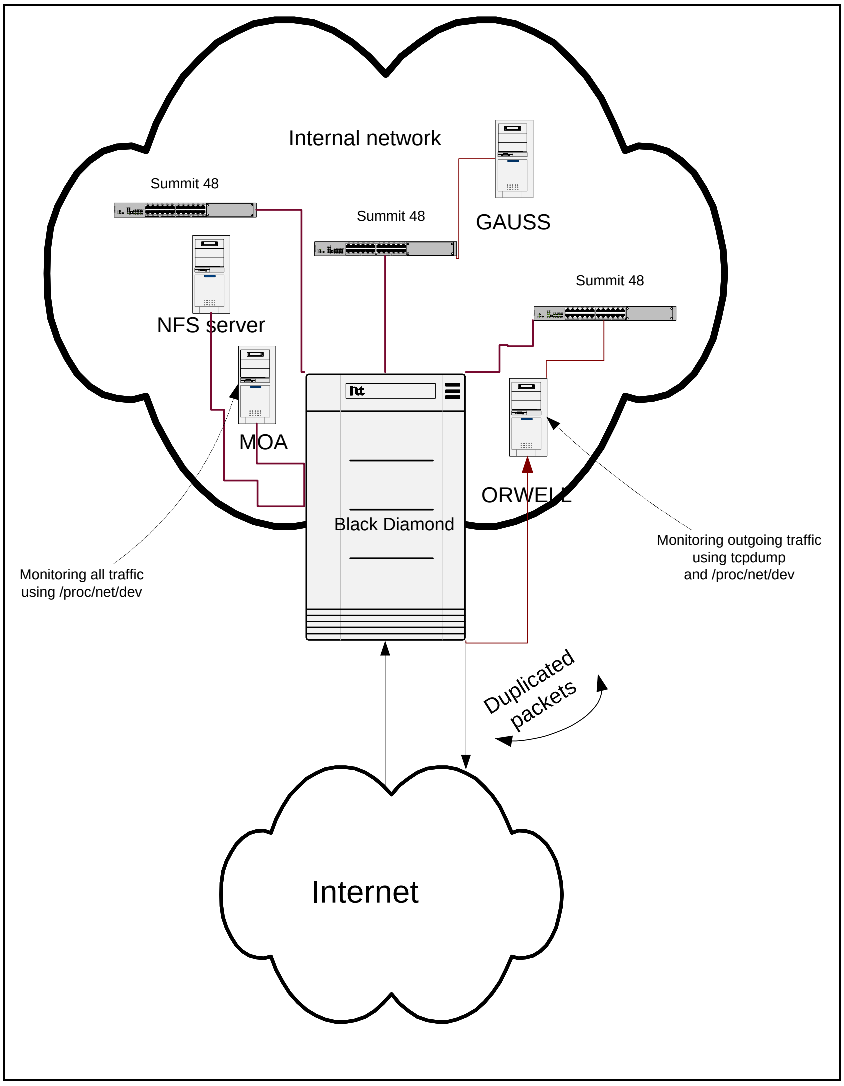

Figure 1: Overview of the departmental network. Note that the activity of the core router due to the internal network is not recorded.The Black Diamond is a router for external traffic and a top-level switch for the internal traffic. Both parts are implemented by different pieces of hardware and are independent of each other.

is similar to the aggregated data of the `tcpdump` measurements although the time resolution is by no means as good. This method has three advantages over `tcpdump`: it usually requires no special operating system privileges, it can be run for a much longer time as only summary data is collected, and it incurs lower intrusion overheads. The tradeoff between resolution and efficiency, both space and time, is addressed later in this paper.

On one of the departmental CPU servers, MOA, we ran the `/proc/net/dev` monitor for twelve days at a measurement interval of one second. The measurement started at 1 February 2002 at 16:24:47.

We have verified the data collected with this method against data collected using `tcpdump` in another context and found both methods to produce the same figures for packet rates.

Whilst network traffic measurements using `tcpdump` provides data which is a time series of a point process measuring entries in the `/proc/net/dev` files provides a counting process with values a_i, where $1 \leq i \leq N$ indicates the time the measurement was taken. The time between measurements is the same period T, $T_N = NT$. Instead of direct measurements the values can be derived by aggregation of measurements. This is illustrated in figure 2. If one assumes

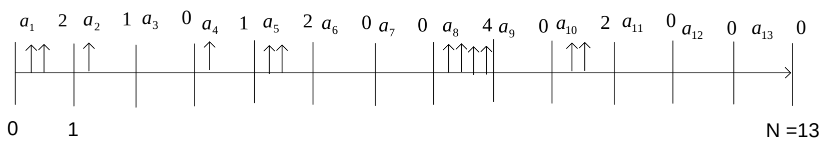

Figure 2: Illustration of aggregating a point process (arrows) to a counting process a_i.

that a time series is generated by a stochastic process, such as Brownian motion, for example, one can investigate the distribution of the *changes* of the a_i from one time step to the next, i.e. the differenced time series

$$\{\Delta a_i = a_{i+1} - a_i | 1 \leq i \leq N-1\}. \tag{1}$$

Alternatively one can work out the distribution of the a_i as they are the changes of the underlying counting process. Depending on the statistical nature, one also expects a different kind of scaling law for these distributions with respect to the length of the aggregation interval T_N.

3 Data analysis

In this section we will show that the probability density functions of the measured file size distributions and the changes in the rates of packets transmitted on the network follow similar power laws. The probability density function $p(x)$ is said to follow a power law if

$$p(x) \propto \beta x^{\gamma}$$

as $x \to \infty$, for $\beta > 0, \gamma < 0$. When investigating the existence or otherwise of a power laws we use exponentially growing bin sizes for the histograms. Apart from the histogram, we also compute the mean and variance of the inter-event times which are useful for distribution fitting.

One of the most famous power laws is Zipf's law, which says that the cumulative probability function P of for instance file size distributions behaves like

$$P(\text{file size} > x) \approx \frac{1}{x} \quad \text{for large } x. \tag{2}$$

One probability density function that can exhibit this behaviour is the Pareto distribution

$$p(x) = \alpha k^{\alpha} x^{-\alpha-1}, \tag{3}$$

where $\alpha, k > 0$ and $x \geq k$. If $\alpha = 1$ the Pareto distribution shows the behaviour of the Zipf law for large x. In a double logarithmic plot, this distribution is a straight line with gradient $-(1+\alpha)$.

3.1 Files and Requests

As described above we have measured distribution on several Linux/UNIX machines. The results are shown in Figure 3. We find that they are extremely well approximated by a Cauchy distribution. The symmetric Cauchy distribution has a pdf given by

$$p(x) = \frac{1}{\pi} \frac{s}{s^2 + x^2}$$

where $s > 0$. In our case we only use the positive half $(x > 0)$of the distribution and therefore have to multiply the equation above by a factor of 2. For large x

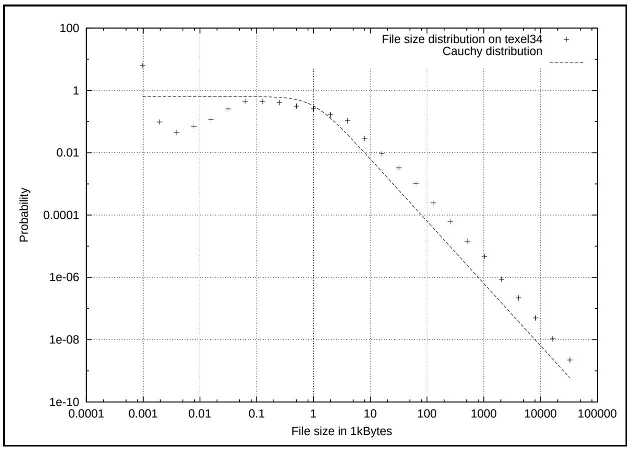

Figure 3: File size distribution on the local disk of a LINUX machine

the distribution behaves like $1/x^2$ so the cumulative distribution function (cdf) therefore obeys Zipf's law.

For the request size distribution we find a very similar distribution, see figure 4. Even thought the actual function looks a bit distorted by some peaks a Cauchy distribution still describes most of the features well. In fact the shape of the curve was hardly changed if we filtered for actual pages, pictures, dynamic or static content.

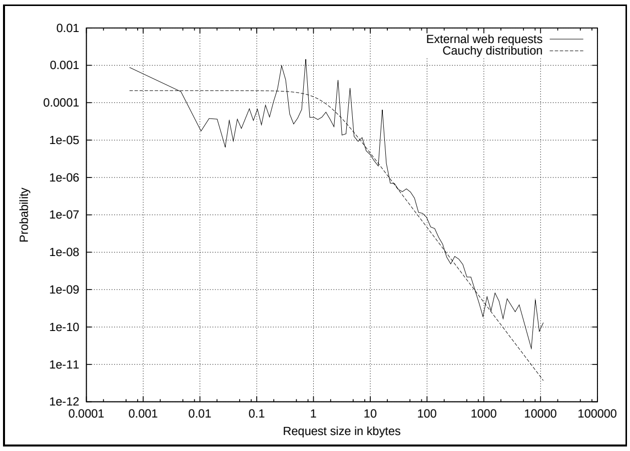

Figure 4: This plot compares the request size distribution of external requests on 22 March 2002 between 12.00 - 14.00 to a Cauchy distribution. Though it is clearly not a perfect fit, it describes the nature of the distribution fairly well, especially the feature of requests below one kilobyte.

The Cauchy distribution is distinguished from the Pareto distribution by its behaviour for small x. This is easily missed when linearly sized bins are used to create the histogram – hence the value of choosing exponentially scaled bin sizes. In [5] the authors use the *log-log complementary distribution* plots to estimate the distribution for request sizes. Our measurements are distinctly non-Pareto but although qualitatively the approximation for small file sizes is not brilliant, it is significantly better than the Pareto assumption as can be clearly seen in figure 4. A Pareto distribution does not describe the request sizes smaller than 1 Kbyte at all.

We also suggest to use a truncated Cauchy distribution [14] whose pdf is given by:

$$p(x) = \begin{cases} \tilde{p}(x)/C & 0 \leq x \leq x_{\max} \\ 0 & \text{else} \end{cases} \tag{4}$$

where C is a normalisation constant

$$C = \int_0^{x_{\max}} \tilde{p}(x) dx.$$

The truncation of the Cauchy distribution gets rid of its usually prohibiting features like infinite moments. The truncation value is naturally given by observed or application imposed maximum file or request sizes.

3.2 CPU server

In this section we present the results for the distribution of packet rates as defined in equation (1). These plots were inspired by work in the analysis of stock prices where changes in prices have been shown to follow similar patterns, see for example [14].

First we look at the entire 12 day measurement period. In a half logarithmic plot, figure 5, we can clearly that the distribution for the network traffic in either direction are clearly not normally distributed as the curves are too leptokurtic. A normal distribution in this plot would resemble a parabola. To investigate

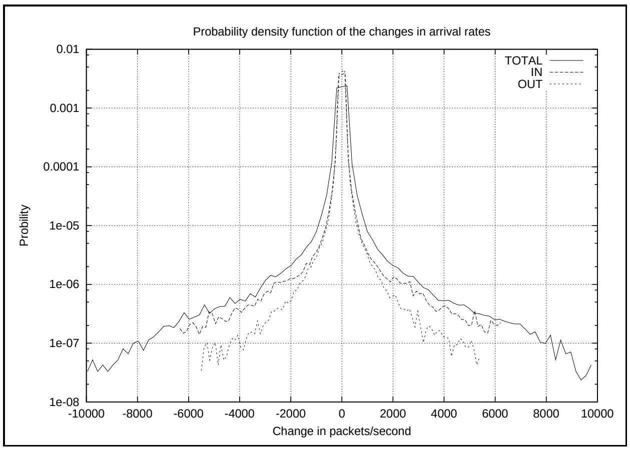

Figure 5: Distribution of the changes in the packet rate over 12 days for the ingoing, outgoing and total traffic to the CPU server MOA in a half logarithmic plot.

the asymptotic behaviour of the changes we re-sampled the distribution and created a double logarithmic plot, figure 6. This time we have added a Cauchy distribution to the plot to substantiate our claim that the distribution well approximated by a Cauchy distribution. As further evidence we have plotted the distribution at different aggregation levels in plot 7. They fall into one master curve when plotted on top of each other as they should do for a Cauchy distribution, see p. 92 in Voit's book and p. 71 in Mantegna's book [14].

We have also restricted the same investigation to the period of the first Monday in the data in the hours between 10am and 5pm. This excludes features the data may show due to daily cycles. We find that there is no change in the nature of features, see figure 8

4 Outlook

In the previous section we have shown that the (truncated) Cauchy distribution can be used to fit both request size and change in packet rate distributions. While this is probably not the best fit, it still makes an analytic description simpler.

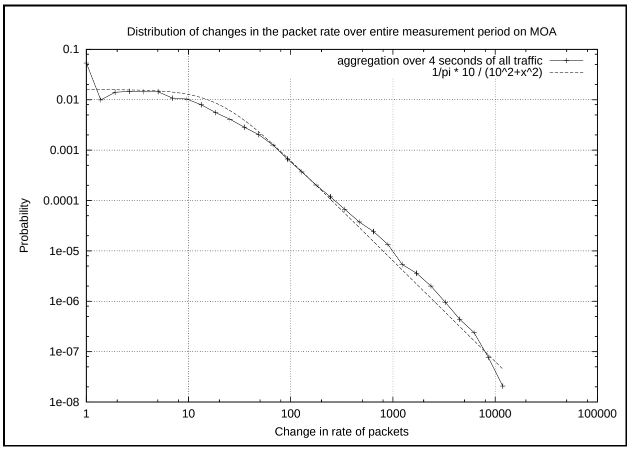

Figure 6: Distribution of the changes in the packet rate over 12 days for the ingoing, outgoing and total traffic to the CPU server MOA in a double logarithmic plot.

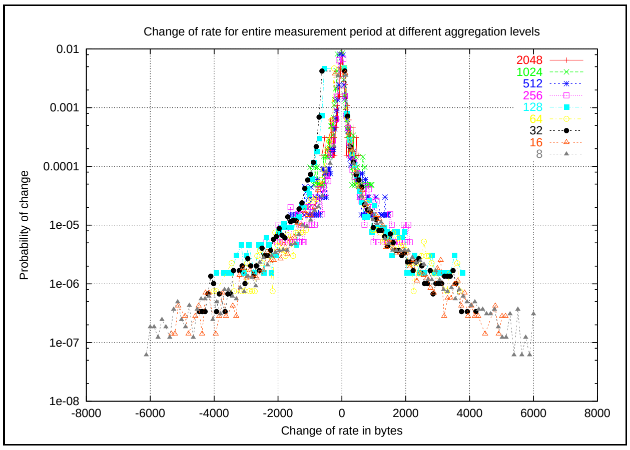

Figure 7: Distribution of the changes in the packet rate over 12 days for the total traffic to the CPU server MOA in a half logarithmic plot for different levels of aggregation.

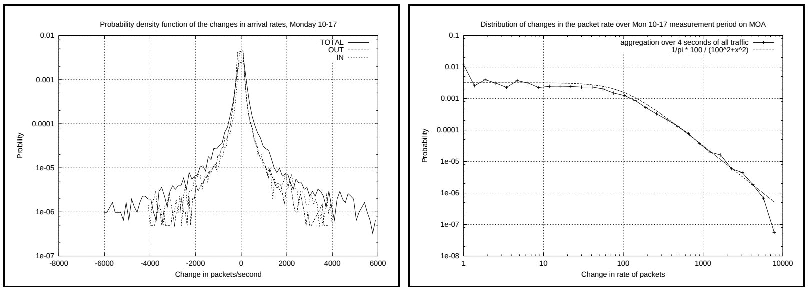

Figure 8: Distribution of the changes in the packet rate on Monday between 10am and 5pm for the ingoing, outgoing and total traffic to the CPU server MOA in a half logarithmic plot (left) and a double logarithmic plot.

The Cauchy, an in fact the normal distribution are special cases stable Lévy distributions. In general Lévy distributions are described by their characteristic function

$$\hat{L}_\mu(z) = \exp(-a|z|^\mu).$$

If the exponent μ is in the range $0 \leq \mu << \leq 2$ the distributions are stable. So, similar to the normal distribution (which is in fact given by $\mu = 2$), they act as fixed points for distributions with infinite variance. The only exponents for which analytic expressions can be written down are the Gaussian and the Cauchy distribution.

Whilst large numbers of Lévy distributions in general do not converge to the normal distribution, truncated Lévy distributions of course do. However, they do it slower than any processes [14].

The fact that we found that a the web requests to the departmental web-server are close to a Cauchy distribution can be used in simulation for instance to make the results more realistic [12]. It is very simple to generate Cauchy distributed random numbers.

Another avenue that we are going to investigate is whether we can use our findings to use an adaptation of the diffusion approximation to calculate server and network loads. In the usual approach the diffusion approximation makes assumption that the queue length at a server can be described by Brownian motion, see for instance [1]. If one changes the normal distribution used for incremental changes in the Brownian motion one gets (truncated) Lévy flights [14]. However, the problem is that these no longer satisfy the diffusion equation which is the essential step in the diffusion approximation. Further research might well show that there is an equivalent equation satisfied by Lévy flights which could be used to continue along the path of the diffusion approximation.

Another possible use of the file and request size distributions we have found may in an extension of the models used in [11]. These models were motivated by the concepts of self-organised criticality [14, 17] which is a theory in Mathematical Physics that is producing numerous simple models that exhibit or explain the abundance of $1/f$ noise found for instance road traffic, sand piles and stock market data.

Acknowledgements

The authors would like to thank Ashok Argent-Katwala, David Thornley and Jeremy Bradley of Imperial College for stimulating discussions.

The research was funded by EPSRC (research grant QUAINT, GR/M80826).

References

[1] L. Kleinrock, *Queueing Systems*, Wiley 1975

[2] Will E. Leland, Murad Taqqu, Walter Willinger and Daniel Wilson, *On the Self-Similar Nature of Ethernet Traffic (Extended Version)*, IEEE/ACM Transactions on Networking, Vol. 2, No. 1, February 1994

[3] Ashok Eramilli, Onuttom Narayan and Walter Willinger, *Experimental Queueing Analysis with Long-Range Dependent Packet Traffic*, IEEE/ACM Transactions on Networking, Vol. 4, No. 2, April 1996

[4] *Unix File Size Survey* `http://www.base.com/gordoni/ufs93.html`

[5] M. Crovella and A. Bestavros, *Self-similarity in world wide web traffic: evidence and possible causes*, IEEE/ACM Transactions on networking, vol. 5, no. 6,(1997) pp. 835-846

[6] Forest Baskett, K. Mani Chandy, Richard R. Muntz and Fernando G. Palacios, *Open, Closed and Mixed Networks of Queues with Different Classes of Customers*, Journal of the ACM 22 (2), pp.248-260, April 1975.

[7] K. Fukuda et al., *Origin of critical behaviour in Ethernet traffic*, `cond-mat/0007435` and references therein

[8] T. Huisinga et al., *A Microscopic Model for Packet Transport in the Internet*, `cond-mat/0102516` and Physica A 294 (2001), pp. 249-256

[9] Andras Veres and Miklos Boda,*The Chaotic Nature of TCP Congestion Control*, INFOCOM (3) 2000, pp. 1715-1723

[10] Hans-Peter Schwefel, *Behavior of TCP-like elastic traffic at a buffered bottleneck router*, `http://citeseer.nj.nec.com/451117.html`; Biplab Sikdar and Kenneth S. Vastola, *The Effect of TCP on the Self-Similarity of Network Traffic*, `http://citeseer.nj.nec.com/sikdar01effect.html`; Allen B. Downey, *The structural cause of file size distributions*, "SIGMETRICS/Performance" 2001, pp. 328-329

[11] D. E. Newman, Nathaniel D. Sizemore, B. A. Carreras and V. E. Lynch, *Growth and Propagation of Disturbances in a Communication Network Model*, Hawaii International Conference on Systems Sciences, January 2002, Hawaii.

[12] Tony Field, Uli Harder and Peter Harrison, *NETWORK TRAFFIC BEHAVIOUR IN SWITCHED ETHERNET SYSTEMS* Technical Report 2002/10 Department of Computing, Imperial College, `http://www.doc.ic.ac.uk/deptechrep/tc0210.html#02-10`

[13] Xiaoyun Zhu, Jie Yu and John Doyle, *Heavy Tails, Generalized Coding and Optimal Web Layout*, IEEE INFOCOM 2001

[14] R. N. Mantegna and H.E. Stanley, *An Introduction to Econophysics*, CUP 2000, ISBN 0 521 62008 2 and J. Voit, *The Statistical Mechanics of Financial Markets*, Springer 2001, ISBN 3 540 41409 6

[15] *Measurement data*, `http://www.doc.ic.ac.uk/~uh/QUAINT/data/`

[16] Tcpdump website `http://www.tcpdump.org/`

[17] P.Bak, C. Tang and K. Wiesenfeld, *Self organised criticality: an explanation of 1/f noise*, Physical Review Letters, *59*, p. 381, (1987); H. J. Jensen, *Self-Organised Criticality*, CUP 1998