

Ambiguities in estimates of critical exponents for long-range dependent processes

Ken Duffy^(1,*), Christopher King⁽²⁾ and David Malone⁽¹⁾

2007

(1) Hamilton Institute, National University of Ireland, Maynooth, Co. Kildare, Ireland

(*)Phone:+353-1-7086463. Fax:+353-1-7086269. E-mail: ken.duffy@nuim.ie

(2) Department of Mathematics, Northeastern University, Boston, MA 02115, USA

Abstract

In this article we show that the Hurst parameter and a well-known wavelet estimator fail to distinguish on/off models with substantially different distributions of on-times. We identify the problem as arising when standard power spectrum techniques, which are well-defined for commonly considered processes, are applied to processes with extremely heavy tailed on-times. We provide an elementary mechanism whereby such extremely heavy-tailed on-time processes arise in a simple queuing model.

PACs: 02.50.-r; 02.50.Ey; 05.40.-a; *MSC:* 62M15.

Keywords: novel long-range dependence; alternating renewal processes; estimation ambiguities.

1 Introduction

Roughly speaking, assuming the process that created a data-set is stationary (ie. has time invariant statistics), the process is said to have long-range dependencies when standard statistical tests on its data-set show correlation structure that decays slower than exponentially as a function of lag. Since the mid-twentieth century, long-range dependence has been observed in data-sets from many distinct fields of physics, engineering and economics (for a few examples see [1, 2, 3, 4, 5, 6, 7, 8])^{*}.

The practical significance of the presence of long-range dependence can not be overestimated, as Long-Range Dependent (LRD) processes exhibit qualitatively different behavior to Short Range Dependent (SRD) processes, motivating developments in both the theory and practice. This is particularly the case in telecommunications, where the physics community has applied its expertise in considering the appearance of LRD processes [11, 12, 13, 14, 15].

From a phenomenological point of view, one important question is: *How can one identify long-range dependence in a process, and how can different LRD processes be reliably distinguished?* We demonstrate that it is possible for standard Hurst-type statistical tests to fail to distinguish different processes

^{*}In some fields it is possible that the apparent presence of long-range dependence can physically be best explained by other phenomena, for example see [9] and [10].

within the class of on/off processes. Processes of this sort arise naturally in many applications areas and thus ambiguities in the class are significant. We describe a one-parameter family of alternating renewal processes with heavy-tailed on-times. In this family of processes we analyse a range of parameter values that has not previously been considered. We distinguish processes with parameters in this extended range by the term “extremely heavy tailed”. For this class we show that the Hurst parameter fails to distinguish between processes whose on-time distribution is heavy tailed and one whose on-time distribution is extremely heavy tailed. We do this by using a regularized power spectrum to compute the Hurst parameter analytically. Although sharing a common Hurst value, the processes for which these ambiguities occur exhibit significantly different qualitative behavior. This suggests that even for on/off processes additional statistical tests are necessary to determine the nature of the source of a data set.

Even with this ambiguity in theory we ask: *would a power spectrum divergence estimator distinguish these processes?* To address this question we estimate the Hurst parameter from simulations using an established wavelet estimator. Both theory and estimation show similar features: the same value of the Hurst parameter for different processes. Finally we introduce a simple queueing model whose output process is extremely heavy tailed. Thus this process would not be distinguished from a process with a more usual heavy tailed on-time distribution.

2 Inferring long-range dependence

The Hurst parameter H , introduced originally in hydrology [1], is a common measure of the long-range dependence of a stochastic process. A Hurst parameter value of greater than 0.5 indicates that the process’s power-spectrum (the Fourier transform of its autocorrelation function) diverges at the origin and signals the presence of long-range dependence. We use the operational approach of taking power-spectrum divergence at the origin as our definition of long-range dependence.

Our first result concerns the LRD “strength” of a particular class of alternating renewal processes whose on-times are distributed as $x^{-\alpha}$ for $0 < \alpha < 2$ and whose off times have finite variance. For $\alpha \in (1, 2]$ we use the standard term “heavy tailed”, as the distribution has finite mean and infinite variance. When $\alpha \in (0, 1]$ we use the term “extremely heavy tailed” as both mean and variance are infinite.

The power-spectrum is not defined for processes with extremely heavy tailed on-times. Processes of this sort are not normally considered in the literature. In order to overcome this difficulty we introduce a regularized version of the power-spectrum and determine its divergence γ at the origin, which is a function of α . The divergence γ is related to the Hurst parameter by $H = (1 + \gamma)/2$. For $1 < \alpha < 2$ the relationship $\gamma = 2 - \alpha$ is deduced in Heath et al. [16]. For $0 < \alpha < 1$ Theorem 1 proves that $\gamma = \alpha$. Thus the Hurst parameter does not differentiate the processes with parameter value $\beta \in (1, 2)$ and $\beta' = 2 - \beta \in (0, 1)$, allowing misclassification unless additional tests are performed.

It is not clear what a power-spectrum estimator would make of a sample from an extremely heavy tailed process. To investigate this question we apply a standard wavelet estimator to a process with $x^{-\alpha}$ on-time distributions. It is demonstrated in Section 4 that the estimator finds the same qualitative behavior as predicted by the theory. That is, a 2 to 1 mapping from parameter values α to divergence γ .

In order to demonstrate that processes with extremely heavy tails can arise in practice we examine

the behavior of on/off processes with short-range dependence that are stalled at a slow roadblock. We model this system by a single server queue fed by an on/off source with short range dependence whose mean arrival-rate is equal to the service-rate. For this system (which we call a balanced queue) we demonstrate the presence of LRD in the departure process assuming infinite waiting space and constant service-rate. This result, stated in Theorem 2, generalizes a well-known result for simple random walks. It proves that the on-times of the departure process from a balanced queue are distributed such that $P(\text{On} > x) \sim x^{-0.5}$. This indicates that Pareto $x^{-\alpha}$ burst-time distributions with $0 < \alpha < 1$ are important in queueing theory. They may appear in systems that includes blocking and delays.

3 The regularized power-spectrum

A process $\{X(t), t \in (-\infty, \infty)\}$ is Wide-Sense Stationary (WSS) if its mean is time invariant, that is $\mathbf{E}[X(t)] = \mathbf{E}[X]$ for all t , and if its autocorrelation function depends only on lag $\mathbf{E}[X(s+t)X(s)] = R_{XX}(t)$. The power-spectrum $S(\theta)$ of a WSS process is the Fourier transform of its autocorrelation function:

$$S(\theta) = \int_{-\infty}^{\infty} e^{i\theta t} R_{XX}(t) dt. \quad (1)$$

Unless $\mathbf{E}[X] = 0$ the expression (1) has an impulse singularity at $\theta = 0$. For this reason it is often preferable to work instead with the Fourier transform of the auto-covariance function $C_{XX}(t) = \mathbf{E}[X(s+t)X(s)] - (\mathbf{E}[X])^2$. The Fourier transform agrees with (1) for $\theta \neq 0$ and is typically continuous at $\theta = 0$. However if the process $\{X(t)\}$ is LRD, then the integrals of both $R_{XX}(t)$ and $C_{XX}(t)$ may be infinite, in which case (1) cannot be used to compute the power-spectrum. We resolve this difficulty by defining a regularized Fourier transform $\tilde{X}_\epsilon(\theta)$ of $X(t)$. We use the expectation of $|\tilde{X}_\epsilon(\theta)|^2$ as an approximation to the power-spectrum.

Assuming that $X(t)$ is uniformly bounded the following integral exists for any $\epsilon > 0$:

$$\tilde{X}_\epsilon(\theta) = \int_{-\infty}^{\infty} X(t) e^{i\theta t} e^{-\epsilon|t|} dt. \quad (2)$$

If $X(t)$ is WSS, then the usual power-spectrum, $S(\theta)$ defined in Equation (1), is recovered using a standard technique [17] from Equation (2) by taking the limit

$$S(\theta) = \lim_{\epsilon \rightarrow 0} \epsilon \mathbf{E} \left[|\tilde{X}_\epsilon(\theta)|^2 \right]. \quad (3)$$

For technical reasons this convergence requires that $X(t)$ be replaced by $X(t) - \mathbf{E}[X]$, in which case the right side of (3) converges to the Fourier transform of the auto-covariance function. As remarked above this distinction is irrelevant for $\theta \neq 0$.

Theorem 1 shows that Equation (2) enables the definition of a regularized power-spectrum for a large class of heavy-tailed on/off processes. The idea is to determine the behavior of $\mathbf{E} \left[|\tilde{X}_\epsilon(\theta)|^2 \right]$, as $\epsilon \rightarrow 0$, and extract its leading order part by re-scaling by ϵ^ν for some $\nu > 0$. This leading order part $S^{\text{reg}}(\theta)$ is defined to be the regularized power-spectrum of the process for $\theta \neq 0$. As Equation (3) shows this procedure yields the correct power-spectrum for a WSS process (with exponent $\nu = 1$). This approach resolves the problem of rescaling with an infinite mean that arises when the standard definition is applied to processes of this sort [12, 18].

Theorem 1 considers alternating renewal processes for which the i.i.d. off periods $\{C_n, n \geq 1\}$ have finite mean and finite variance. The on i.i.d. periods $\{B_n, n \geq 1\}$ may be heavy tailed, as occurs for the output of the balanced queue. The process $X(t)$ is defined to be 0 for $t < 0$ and for $t \geq 0$ it is defined as follows:

$$X(t) = \begin{cases} 1 & \text{if } t \in [\sum_{i=1}^n (B_i + C_i), \sum_{i=1}^n (B_i + C_i) + B_{n+1}) \text{ for some } n \geq 0 \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

with the empty sum $\sum_{i=1}^0 (B_i + C_i)$ defined to be zero. Without loss of generality the rate is normalized to be 1 during an on period. There is no assumption of stationarity for the process, which always starts with an on period at $t = 0$.

For the purposes of Theorem 1 it is convenient to describe the heavy-tailed distribution of the on periods by specifying the asymptotic behavior of their density function. The duration of an on period is denoted B and its density denoted f_B . Hence, for example,

$$P(B > t) = \int_t^\infty f_B(u) du. \quad (5)$$

For some $\alpha > 0$, f_B is assumed to satisfy

$$\lim_{u \rightarrow \infty} u^{\alpha+1} f_B(u) = c > 0 \quad (6)$$

or, more compactly, $f_B(u) \sim u^{-\alpha-1}$. Using (5) this implies that $P(B > t) \sim t^{-\alpha}$.

For $\alpha > 2$ both the mean and variance are finite, thus B is short range. For $1 < \alpha \leq 2$, the mean is finite but the variance is infinite. For $0 < \alpha \leq 1$ both the mean and variance are infinite.

Theorem 1 *Let $X(t)$ be the alternating renewal process defined in (4) and define ν by:*

$$\nu = \begin{cases} 1 & \text{for } \alpha > 1; \\ \alpha & \text{for } 0 < \alpha < 1. \end{cases}$$

Then for all $\theta \neq 0$ the following limit exists and is not identically zero:

$$S^{\text{reg}}(\theta) = \lim_{\epsilon \rightarrow 0} \epsilon^\nu \mathbf{E} \left[|\tilde{X}_\epsilon(\theta)|^2 \right]. \quad (7)$$

Furthermore the following asymptotic behavior holds:

$$S^{\text{reg}}(\theta) \sim \theta^{-\gamma} \quad \text{as } \theta \rightarrow 0, \quad (8)$$

where

$$\gamma = \begin{cases} 0 & \text{for } \alpha > 2, \\ 2 - \alpha & \text{for } 1 < \alpha < 2, \\ \alpha & \text{for } 0 < \alpha < 1. \end{cases}$$

The proof appears in Appendix A.

A graph of γ vs. α is shown in Figure 1. The shape of this graph for $\alpha > 1$ can (almost) be deduced from Theorem 4.3 of [16] and Proposition 4 of [19]. One of the main contributions of this paper is to determine the relationship for $0 < \alpha < 1$. The exponent γ is related to the Hurst parameter by $H = (1 + \gamma)/2$. Thus Theorem 1 demonstrates that the same Hurst parameter arises for on/off sources with completely different heavy-tailed on distributions.




Figure 1: Power-spectrum divergence, γ , against on period power-tail strength, α .

4 Application of a standard power-spectrum estimator

The purpose of the experiments in this section is to answer the following question: *if presented with a sample of data from a process described in the previous section, what would the application of a standard power-spectrum estimator tell you?* Given a sample of data of unknown origin and assuming its source to be WSS it is likely such an estimator would be used.

Many methodologies [20] can be employed to determine the “strength” of the long-range dependence of the source of a time-series. For a comparison of a selection of those invented prior to 1995 see [21]. One of the most successful approaches, and the one that we use, is described in [22]. It consists of a wavelet-based estimator for the power-spectrum defined in Equation (1). This enables one to estimate the divergence γ of $S(\theta)$ at the origin of a WSS source by the following procedure. For a range of values of n plot $\log(S(1/2^n))$ against n . If the graph is of constant slope going to $-\infty$, then the time-series is constant. If the graph is of positive slope, then the value of the slope is the estimate of γ . If the graph raises steeply and then levels to a flat line, the conclusion is that the source of the time-series is short-range dependent.

To test the application of the wavelet-based estimator a variety of on/off sources as described in Equation (4) were simulated. These had exponentially distributed off periods and on periods distributed as $P(B > x) \sim x^{-\alpha}$, for a range of values of $\alpha \in (0, 2]$. Veitch [23] has made available `matlab` code for a wavelet-based power-spectrum estimator. We do not use that implementation in this section as the infinite expected burst length associated with $0 < \alpha \leq 1$ makes it impractical. Instead an implementation was designed to exploit the known response of the wavelet estimator to long, constant bursts. C-code for the burst-based estimator can be found online [24].

In these simulations the probability an off period is longer than $l \geq 1$ is 2^{1-l} . On periods last for $u^{-1/\alpha}$ where u is uniformly distributed in $(0, 1]$. Each source comprises 1,000,000 on/off cycles. The graph on the left of Figure 2 consists of log-log plots produced by the wavelet estimator for a range of values of α . Note the graphs form straight lines. Thus the conclusion from this test is that the source exhibits long-range dependence.

The divergence of the power-spectrum at the origin is measured by the slope of these lines. The graph

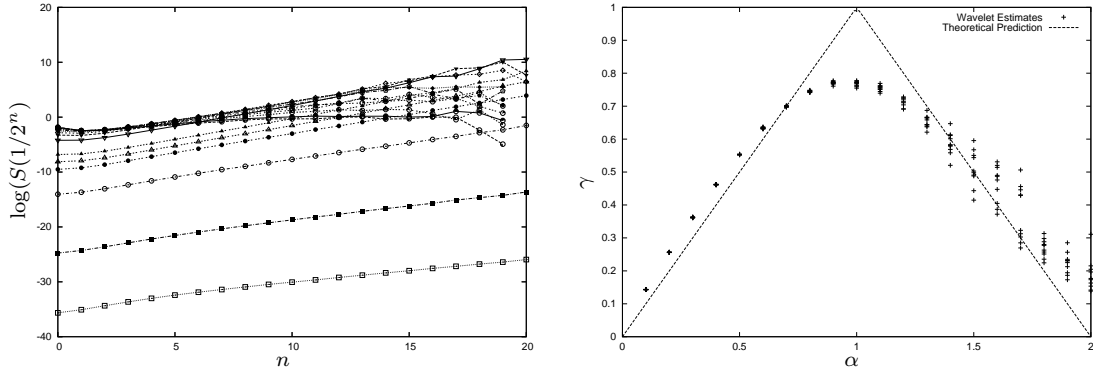


Figure 2: Wavelet estimated power spectrum for $0 < \alpha < 2$.




Figure 3: Empirical power spectrum and estimated divergence $\alpha = 0.4$.

on the right of Figure 2 shows least-squares fitted slopes, fit between octaves 3 and 15, for different values of α . Ten sources with different seeds are presented for each α . Theorem 1 predicts the hat function, which is shown for comparison.

For many values of γ it is clear that the estimator does not distinguish between two sources, one with $\alpha \in (1, 2)$ and one with $\alpha \in (0, 1)$. Thus, in the absence of additional tests, incorrect inferences and implications could be drawn. For illustration the graph in Figure 3 shows the estimated power spectrum $S(\theta)$ vs θ for $\alpha = 0.4$ and the estimated divergence $\theta^{-0.46226}$ given by the wavelet estimator described above.

As suggested by an anonymous reviewer, to illustrate that the agreement of theory and estimate is not a peculiarity of the wavelet approach, we implemented another estimator. We selected the variance-time plot method. This works by estimating the slope of the log sample variance of the aggregated process, $X^m(t) = \int_t^{t+m} X(s)ds/m$, versus log the aggregation factor, m ; details can be found in [20]. Due to numerical limitations, it was not possible for us to treat values of $\alpha < 0.5$, but Figure 4 shows excellent agreement with theory.




Figure 4: Variance-time plot estimates for $0.5 \leq \alpha < 2$.

5 A balanced queue

In this section we demonstrate how the mechanics of a simple physical system can change a short range dependent process into an extremely heavy tailed process. The model is a single server queue with infinite waiting space whose service rate equals the mean arrival rate.

The balanced queue is analyzed using a continuous-time fluid model. In this model a single server is fed by an on/off arrival process that carries traffic at constant rate r during on periods and at rate 0 during off periods. The durations of on periods $\{A_n\}$ are assumed to be drawn independently from a distribution with finite mean $\mathbf{E}[A]$ and variance. The durations of the off periods are assumed to be drawn independently from an exponential distribution of rate λ .

The server has an infinite buffer and serves at constant rate s whenever its buffer is non-empty. We assume that $r > s$ so that the server is idle only when both the buffer is empty and the arrival-rate is zero. This ensures that the departure process forms an alternating sequence of rate- s on periods and rate-0 off periods. As the off periods of the arrival process are exponentially distributed, the on and off periods of the departure process are independent and the off periods are i.i.d. exponential with rate λ . Our goal is to analyse the distribution of the on periods of this departure process.

This system has been analyzed in detail by other authors. A thorough review can be found in Boxma and Dumas [25]. They define a parameter called the traffic intensity

$$\rho = \frac{r\lambda\mathbf{E}[A]}{s(1 + \lambda\mathbf{E}[A])}$$

that is the ratio of the arrival rate of traffic to the maximum possible service rate. The condition for stability is $\rho \leq 1$. We define the *balanced queue* by the condition $\rho = 1$, which is the borderline case for stability.

Theorem 2 is a result regarding the output process from a balanced queue. It is assumed that the arrival process starts at time $t = 0$ with an on period. Hence the departure process is not stationary. Indeed in some cases the departure process cannot converge to a stationary distribution.

Theorem 2 *Let $\{B_1, B_2, \dots\}$ be the sequence of on period durations of the departure process from the model described above. Assuming $\rho = 1$, the durations $\{B_n\}$ are i.i.d. and their distribution, B ,*

satisfies

$$\mathbf{P}(B > x) \sim x^{-1/2} \quad \text{as } x \rightarrow \infty.$$

The proof appears in Appendix A.

Theorem 2 shows that the output from the balanced queue fed with a SRD source has power-tailed on periods with $\alpha = 0.5^\dagger$.

6 Conclusions

We demonstrate that the Hurst parameter and a wavelet estimator cannot reliably distinguish different on/off LRD processes. We do this by considering a class of processes not previously treated in the literature. This suggests that one must perform additional statistical tests to characterize the source of long-range dependences or risk drawing incorrect inferences. We illustrate that processes from the class considered can be created by the mechanics of a simple physical system.

Acknowledgment: The work of K.D and D.M. is supported by Science Foundation Ireland grant 03/IN3/I396. C.K.'s work is supported in part by National Science Foundation Grant DMS-0400426.

A Proofs

As well as results from [25] the proof uses the following Tauberian theorem regarding the tail of a probability distribution and its Laplace transform from Feller [27] p.447:

Lemma 1 *Let $F(x) = \mathbf{P}(X \leq x)$ be the distribution of a positive random variable X and let $\phi(t) = \mathbf{E}[e^{-tX}]$ be its Laplace transform. Then each of the following statements implies the other:*

$$\begin{aligned} 1 - \phi(t) &\sim t^{1-a} L\left(\frac{1}{t}\right) \quad \text{as } t \rightarrow 0; \\ 1 - F(x) &\sim \frac{1}{\Gamma(a)} x^{a-1} L(x) \quad \text{as } x \rightarrow \infty; \end{aligned}$$

where $a \geq 0$ and $L(x)$ is a slowly varying function. That is, L is a measurable function with the asymptotic behavior

$$\frac{L(tx)}{L(t)} \rightarrow 1 \quad \text{as } t \rightarrow \infty, \quad \text{for every fixed } x.$$

Proof of Theorem 1: the on/off process is simple enough to allow an explicit calculation of its regularized Fourier transform defined in Equation (2). For $k = 1, 2, \dots$, let B_k denote the duration of the k^{th} busy period. Suppose that the k^{th} period starts at time σ_k and ends at time τ_k . From the definition (4) it follows that $\sigma_1 = 0$. Defining

$$z = i\theta - \epsilon, \tag{9}$$

[†]For another interesting result on balanced queues see [26] where it is proved that the busy period tail in the case of regular variation with $\rho = 1$ behaves like $x^{-1/\nu}$ if the service time tail is like $x^{-\nu}$.

it follows that

$$\tilde{X}_\epsilon(\theta) = \frac{1}{z} \sum_{k=1}^{\infty} (e^{z\tau_k} - e^{z\sigma_k}).$$

Hence, the regularized power is computed using

$$\mathbf{E}[|\tilde{X}_\epsilon(\theta)|^2] = \frac{1}{|z|^2} \sum_{k,l=1}^{\infty} \mathbf{E}[(e^{\bar{z}\tau_k} - e^{\bar{z}\sigma_k})(e^{z\tau_l} - e^{z\sigma_l})].$$

As busy and idle periods of the process are independent, this expectation can be computed in terms of the following generating functions for a typical on period B and off period C :

$$\phi(w) = \mathbf{E}[e^{wB}] \quad \text{and} \quad \psi(w) = \mathbf{E}[e^{wC}],$$

where these are well-defined when w has a non-positive real part. Recalling Equation (9), the result is

$$\begin{aligned} \mathbf{E}[|\tilde{X}_\epsilon(\theta)|^2] &= \frac{\phi(z + \bar{z}) - \phi(z) - \phi(\bar{z}) + 1}{|z|^2(1 - \phi(z + \bar{z})\psi(z + \bar{z}))} \\ &\quad + \frac{(\phi(\bar{z}) - 1)(\phi(z + \bar{z}) - \phi(\bar{z}))\psi(\bar{z})}{|z|^2(1 - \phi(\bar{z})\psi(\bar{z}))(1 - \phi(z + \bar{z})\psi(z + \bar{z}))} \\ &\quad + \frac{(\phi(z) - 1)(\phi(z + \bar{z}) - \phi(z))\psi(z)}{|z|^2(1 - \phi(z)\psi(z))(1 - \phi(z + \bar{z})\psi(z + \bar{z}))}. \end{aligned} \tag{10}$$

In order to compute $S^{\text{reg}}(\theta)$, it is necessary to extract the leading order part from Equation (10) as $\epsilon \rightarrow 0$. For $\theta \neq 0$, the only singularities on the right side of (10) arise from the term $(1 - \phi(z + \bar{z})\psi(z + \bar{z}))$ in the denominators, since $z \rightarrow i\theta$ as $\epsilon \rightarrow 0$, implies that $z + \bar{z} \rightarrow 0$.

The behavior of this term depends on the tail behavior of the on period, B . If $\alpha > 1$ in Equation (6), then $\mathbf{E}[B] < \infty$, which implies that

$$1 - \phi(z + \bar{z})\psi(z + \bar{z}) \sim 2\epsilon(\mathbf{E}[B] + \mathbf{E}[C]).$$

Therefore, to leading order in this case, (10) behaves as

$$\begin{aligned} \mathbf{E}[|\tilde{X}_\epsilon(\theta)|^2] &\sim \frac{1}{2\theta^2\epsilon(\mathbf{E}[B] + \mathbf{E}[C])} \left[2 - \phi(i\theta) - \phi(-i\theta) \right. \\ &\quad + \frac{(\phi(-i\theta) - 1)(1 - \phi(-i\theta))\psi(-i\theta)}{(1 - \phi(-i\theta)\psi(-i\theta))} \\ &\quad \left. + \frac{(\phi(i\theta) - 1)(1 - \phi(i\theta))\psi(i\theta)}{(1 - \phi(i\theta)\psi(i\theta))} \right]. \end{aligned} \tag{11}$$

The right side diverges as ϵ^{-1} and this establishes (7) for $\alpha > 1$.

For $\alpha \leq 1$, the mean of B is infinite and more care is needed. Defining g by

$$\mathbf{E}[e^{-2\epsilon B}] = 1 - g(\epsilon),$$

the leading order part of (10) becomes

$$\begin{aligned} \mathbf{E} \left[|\tilde{X}_\epsilon(\theta)|^2 \right] &\sim \frac{1}{\theta^2(2\epsilon\mathbf{E}[C] + g(\epsilon))} \left[2 - \phi(i\theta) - \phi(-i\theta) \right. \\ &\quad + \frac{(\phi(-i\theta) - 1)(1 - \phi(-i\theta))\psi(-i\theta)}{(1 - \phi(-i\theta)\psi(-i\theta))} \\ &\quad \left. + \frac{(\phi(i\theta) - 1)(1 - \phi(i\theta))\psi(i\theta)}{(1 - \phi(i\theta)\psi(i\theta))} \right]. \end{aligned} \quad (12)$$

Applying Lemma 1 with $\alpha \leq 1$ gives $g(\epsilon) \sim \epsilon^\alpha$ as $\epsilon \rightarrow 0$. Hence, $2\epsilon\mathbf{E}[C] + g(\epsilon) \sim \epsilon^\alpha$ and, thus, Equation (7) follows for this case.

The behavior of $S^{\text{reg}}(\theta)$ as $\theta \rightarrow 0$ is determined in a similar way. If $\alpha > 2$, both the on and off periods have finite mean and variance. Evaluating the expression (11) at $\theta = 0$ yields a finite constant, establishing (8) for this case.

For $1 < \alpha < 2$, defining

$$h(\theta) = \phi(i\theta) - 1 - i\theta\mathbf{E}[B] = \mathbf{E}[e^{i\theta B} - 1 - i\theta B],$$

a similar analysis shows that the leading order part of $S^{\text{reg}}(\theta)$ as $\theta \rightarrow 0$ goes as $\text{Re}(h(\theta))/\theta^2$. To determine the asymptotic behavior of $h(\theta)$ as $\theta \rightarrow 0$, write

$$\begin{aligned} h(\theta) &= \int_0^\infty (e^{i\theta x} - 1 - i\theta x) f_B(x) dx \\ &= \int_0^\infty \left(\frac{e^{i\theta x} - 1 - i\theta x}{x^{\alpha+1}} \right) x^{\alpha+1} f_B(x) dx \\ &= \theta^\alpha \int_0^\infty \left(\frac{e^{iu} - 1 - iu}{u^{\alpha+1}} \right) \left(\frac{u}{\theta} \right)^{\alpha+1} f_B\left(\frac{u}{\theta}\right) du. \end{aligned}$$

Since, for $1 < \alpha < 2$, the function $(e^{iu} - 1 - iu)/u^{\alpha+1}$ is integrable on $[0, \infty)$, the dominated convergence theorem and (6) imply

$$\lim_{\theta \rightarrow 0} \theta^{-\alpha} h(\theta) = c \int_0^\infty \left(\frac{e^{iu} - 1 - iu}{u^{\alpha+1}} \right) du,$$

establishing (8) for $1 < \alpha < 2$.

For the remaining case, $0 < \alpha < 1$, define

$$q(\theta) = 1 - \phi(i\theta), \quad r(\theta) = 1 - \psi(i\theta) \quad \text{and} \quad k = \lim_{\epsilon \rightarrow 0} \epsilon^{-\alpha} g(\epsilon).$$

Renormalising and rewriting Equation (12) gives

$$\frac{1}{\theta^2 k} \text{Re} \left[\frac{q(\theta)r(\theta)}{r(\theta) + (1 - r(\theta))q(\theta)} \right] \quad \text{as } \epsilon \rightarrow 0. \quad (13)$$

Since, as $\theta \rightarrow 0$, $q(\theta) \sim \theta^\alpha$ and $r(\theta) \sim i\theta$, it follows that the leading order part of (13) is

$$\frac{1}{\theta^2 k} \text{Re} \left[- \frac{r(\theta)^2}{q(\theta)} \right],$$

which establishes (8) for $0 < \alpha < 1$.

QED

Proof of Theorem 2: Recall that A and B are the durations of on periods of the arrival and departure processes respectively. Define their Laplace transforms

$$\gamma(u) = \mathbf{E}[e^{-uA}] \quad \text{and} \quad \pi(u) = \mathbf{E}[e^{-uB}].$$

Let $r' = r/s$, then the following formula is derived in Theorem 3.6 of [25] (though in different notation):

$$\pi(u) = \gamma\left(r'u + \lambda(r' - 1)(1 - \pi(u))\right). \quad (14)$$

Defining

$$\kappa = \mathbf{E}[A] \sqrt{\frac{2(\mathbf{E}[A] + 1/\lambda)}{\mathbf{E}[A^2]}},$$

a straightforward calculation using (14) shows that

$$\lim_{u \rightarrow 0^+} \left(\frac{1 - \pi(u)}{\sqrt{u}} \right) = \kappa.$$

Applying Lemma 1 with $a = 1/2$ and $L = \kappa$ gives the result:

$$\lim_{x \rightarrow \infty} x^{1/2} \mathbf{P}(B > x) = \frac{2\kappa}{\sqrt{\pi}}.$$

QED

References

- [1] H. E. Hurst. Long-term storage capacity of reservoirs. *Transactions of the American Society of Civil Engineers*, 116:770–808, 1951.
- [2] B. B. Mandelbrot and J. Wallis. Some long-run properties of geophysical records. *Water Resources Research*, 5:321–340, 1969.
- [3] T. Musha and H. Higuchi. The 1/f fluctuation of a traffic current on an expressway. *Jpn. J. Appl. Phys.*, 15(7):1271–1275, 1976.
- [4] M. Agu, H. Akabane, and T. Saito. The 1/f fluctuation of the flow of particles moving in a viscous fluid. *Jpn. J. Appl. Phys.*, 70(9):2798–2799, 2001.
- [5] S. V. Dudley, A. L. Goldberger, S. Havlin, R. N. Mantegna, M. E. Matsa, M. Simons C.-K. Peng, and H. E. Stanley. Long range correlation properties of coding and noncoding DNA sequences: Genbank analysis. *Physical Review E*, 51:5084–5091, 1995.
- [6] M. E. Crovella and A. Bestavros. Self-similarity in world wide web traffic: Evidence and possible causes. *IEEE/ACM Transactions on Networking*, 5(6):835–846, 1997.

- [7] N. J. Cutland, P.E. Kopp, and W. Willinger. Stock price returns and the Joseph effect. *Progress in Probability*, 37:327–351, 1995.
- [8] C. M. Hurvich and P. Soulier. Testing for long memory in volatility. *Econometric theory*, 18:1291–1308, 2002.
- [9] V. Klèmes. The Hurst phenomenon: a puzzle? *Water Resources Research*, 10:675–688, 1974.
- [10] R. Oliver and J.L. Ballester. Is there memory in solar activity? *Physical Review E*, 58(5):5650–5654, 1998.
- [11] I. Casabi. 1/f noise in computer network traffic. *J. Phys. A*, 27:417–421, 1994.
- [12] M. Takayasu, H. Takayasu, and T. Sato. Critical behaviors and 1/f noise in information traffic. *Physica A*, 233:824–834, 1996.
- [13] G. Caldarelli, R. Marchetti, and L. Pietronero. The fractal properties of Internet. *Europhysics Letters*, 52:386–391, 2000.
- [14] M. Takayasu, H. Takayasu, and K. Fukuda. Dynamic phase transition observed in the Internet traffic flow. *Physica A*, 277:248–255, 2000.
- [15] K. Fukuda, M. Takayasu, and H. Takayasu. Origin of critical behavior in Ethernet traffic. *Physica A*, 287:289–301, 2000.
- [16] D. Heath, S. I. Resnick, and G. Samorodnitsky. Heavy tails and long range dependence in on/off processes and associated fluid models. *Mathematics of Operations Research*, 23(1):145–165, 1998.
- [17] Alberto Leon-Garcia. *Probability and Random Processes for Electrical Engineering*. Addison Wesley, 1989.
- [18] M. Takayasu, A. Yu, K. Fukuda, and H. Takayasu. *Phase transition and 1/f noise in the Internet Packet Transport*, pages 57–74. Traffic and Granular Flow 97. Springer, Singapore, 1998.
- [19] J. A. Gubner. Theorems and fallacies in the theory of long-range-dependent processes. *IEEE Transactions on Information Theory*, 51(3):1234–1239, 2005.
- [20] J. Beran. *Statistics for long-memory processes*. Chapman & Hall, New York, 1994.
- [21] M. S. Taqqu, V. Teverovsky, and W. Willinger. Estimators for long-range dependence: an empirical study. *Fractals*, 3(4):785–798, 1995.
- [22] P. Abry and D. Veitch. Wavelet analysis of long-range-dependent traffic. *IEEE Transactions on Information Theory*, 44(1):2–15, 1998.
- [23] D. Veitch. Matlab wavelet code. Available: http://www.emulab.ee.mu.oz.au/~darryl/secondorder_code.html.
- [24] D. Malone. Wavelet C-code. Available: <http://www.hamilton.ie/dwmalone/wavelet/>.
- [25] O. J. Boxma and V. Dumas. Fluid queues with long-tailed activity period distributions. *Computer Communications*, 21(17):1509–1529, 1998.

- [26] A. P. Zwart. Tail asymptotics for the busy period in the GI/G/1 queue. *Mathematics of Operations Research*, 26:485–493, 2001.
- [27] W. Feller. *An introduction to probability theory and its applications*, volume II. Wiley and Sons, New York, 1971.