

Heavy Tails, Generalized Coding, and Optimal Web Layout

Xiaoyun Zhu Jie Yu John Doyle
xyzhu, jyu, doyle@cds.caltech.edu
Control & Dynamical Systems 107-81
California Institute of Technology
Pasadena, CA 91125

Abstract—This paper considers Web layout design in the spirit of source coding for data compression and rate distortion theory, with the aim of minimizing the average size of files downloaded during Web browsing sessions. The novel aspect here is that the object of design is layout rather than code-word selection, and is subject to navigability constraints. This produces statistics for file transfers that are heavy tailed, completely unlike standard Shannon theory, and provides a natural and plausible explanation for the origin of observed power laws in Web traffic. We introduce a series of theoretical and simulation models for optimal Web layout design with varying levels of analytic tractability and realism with respect to modeling of structure, hyperlinks, and user behavior. All models produce power laws which are striking both for their consistency with each other and with observed data, and their robustness to modeling assumptions. These results suggest that heavy tails are a permanent and ubiquitous feature of Internet traffic, and not an artifice of current applications or user behavior. They also suggest new ways of thinking about protocol design that combines insights from information and control theory with traditional networking.

I. INTRODUCTION

STATISTICAL analysis on empirical data has shown strong indication of the *self-similar* nature of the traffic in both local area networks (LANs) [1] and wide area networks (WANs) [2], quite unlike the traditionally assumed Poisson traffic models. Real traffic exhibits long-range dependence and high burstiness over a wide range of time scales, in contrast with Poisson model's characteristic burst length that is smoothed out by averaging [3]. These discoveries have inspired extensive research in the modeling of network traffic statistics, their relationship to network protocols and the impact on network performance. For example, analysis and simulation results have shown a significant difference in queuing performance between traditional traffic models and self-similar models [4].

A detailed understanding of the origins of self-similar network traffic has begun to emerge. It has been widely argued that the dominant effect is due to heavy-tailed Web traffic being streamed out onto the network to create long-range correlations in packet rates. An important class of *heavy-tailed distributions* have power law tails, i.e., $P[X > x] \sim cx^{-\alpha}$, as $x \rightarrow \infty$, $0 < \alpha < 2$, $c > 0$. In [5] it was shown that the aggregate flows of many ON/OFF processes, each with power law tails with exponent α in their time distributions produce long-range correlations with Hurst parameter $H = (3 - \alpha)/2$. Measurements confirmed that the distribution of the transmission times or sizes of individual files does indeed have a power law tail with $\alpha \approx 1$ in 1995 [6] and $\alpha \approx 1.4$ in 1998 [7]. The resulting H also matched the predictions in [5]. We will neglect the difference between the file sizes and the corresponding transmission times and use the term “Web file transfers” to refer to both.

Given the notion that self-similarity in network traffic can be traced to heavy-tailed Web file transfers, a natural next question is what is the origin of heavy tails in Web files, and in partic-

ular, why are the statistics so consistent with a power law with $\alpha \approx 1$ or slightly bigger. Is this due to some anomaly that might disappear as Web layout design evolves, or as new applications replace the Web as the dominant source of Internet traffic, or is it an intrinsic, natural, and permanent feature that can be expected to persist? This paper will argue strongly for the latter position, expanding on the theme of Highly Optimized Tolerance (HOT) [8], [9], [10], [11] that claims that power laws are the ubiquitous outcome of robust design of complex systems. HOT systems arise when deliberate robust design aims for a specific level of *tolerance* to uncertainty, which is traded off against the cost of the compensating resources. *Optimization* of this trade-off may be associated with some mixture of explicit planning as in engineering, or mutation and natural selection as in biology, but the word “design” is used loosely to encompass both. The claim that power laws are ubiquitous in natural and man-made systems is not new, nor is it without controversy. The importance of power laws has been particularly promoted by advocates of self-organized criticality (SOC) [12]. HOT offers a radically different alternative theory for the nature of complexity and the origin of power laws and “phase transitions” in complex systems.

While there have been recent network-oriented extensions of Shannon theory for both source [13] and channel [14] coding, many central problems related to coding in networks have received essentially no theoretical treatment. For example, if Web sites and the clients browsing them are viewed collectively as a single aggregate “source,” then this source involves both feedback and geometry as users interactively navigate context sensitive hyperlinked content. The “channel” losses experienced by this source are primarily due to congestion caused by traffic generated by the source itself. While coding theory is relevant to file compression, and is widely used, the geometric and dynamic feedback aspects are less familiar. Furthermore, the full implications of the research on heavy-tailed Web and network traffic have yet to be understood or exploited. In short, there has been no coherent coding or control theoretic treatment of the geometric and dynamic aspects of Web and other application traffic. Our central aim is to begin such treatment of Web layout and access in the spirit of standard data compression and rate distortion theory. Since a natural and direct treatment of the problem does not appear to be possible using existing techniques, an entirely new “generalized source coding” theory must be developed. This generalization does have the appealing pedagogical feature of retaining the standard theory as a special case, however.

The connection between HOT and WWW files was first discussed in [8], and the connections with source coding were first

made explicit in [10], which also showed that HOT explained forest fire data much better than the standard SOC results. We will quickly review this WWW work, which provides insights into the origin of power law distributions but is highly abstracted in order to obtain analytic solutions. The bulk of this paper then explores how the assumptions in these models can be relaxed in various ways with some loss in the transparency of the solutions, and discusses in more detail how realistic the assumptions are, and how much the results tell us about real Web layout. We will assume the reader is familiar with standard coding and rate distortion theory, but will review certain aspects that are particularly related to our generalization.

II. SOURCE CODING AND RATE DISTORTION THEORY

The general problem of optimally compressing files is well-known to be undecidable, and one of the brilliant insights of Shannon theory is to focus instead on a more computationally tractable *stochastic* relaxation. Files to be compressed are replaced by ensembles of source symbols, with probabilities replacing frequencies of occurrence in a file. The relaxed problem turns out to be much easier than the former, and its solution can be shown to be provably close to those of the original one. Thus a source message is assumed to be a sequence of independent, identically distributed (IID) random variables, chosen from N source symbols, which occur with probabilities p_i , $\sum_i p_i = 1$. Source coding yields a map that assigns each source symbol a codeword c_i of length l_i in a D -ary alphabet $\{0, 1, \dots, D-1\}$, subject to Kraft's inequality constraint $\sum_i D^{-l_i} \leq 1$, which is equivalent to instantaneous decodability. Defining a resource $r_i = D^{-l_i}$ (an unconventional notation that will be useful later), optimal lossless compression minimizes the average coded file length subject to decodability, or

$$J = \left\{ \sum p_i l_i \mid l_i = f(r_i), \sum r_i \leq R \right\}. \quad (1)$$

with $f(r) = -c \log(r)$, $c = 1/\log(D)$ and $R = 1$. The optimal cost $J_0 = -\sum p_i \log(p_i)$ is the Shannon-Kolmogorov entropy. We will refer to this formulation as the PLR model, for the three elements of probabilities, lengths, and resources.

To connect and contrast Web layout with standard rate distortion techniques, suppose for concreteness that we are interested in a Web site that would be used to browse, search and locate (large) high resolution photographs or other images of interest from a large data base. We'll discuss other types of Web and other application media later. Typically Web sites for such an application will create a variety of pages specifically to help the user navigate the Web site, such as thumbnails of low resolution grouped by topics or features. It is worth reviewing exactly what problem rate distortion theory addresses that is relevant to this problem, as it involves lossy coding. Conventional rate distortion methods can be used to convert a single high resolution image into a sequence of lower resolution images that provide a tradeoff between compressed file size and distortion, and these can be hyperlinked from low to medium to high resolution, so that a user can progressively obtain higher resolution at the expense of larger file downloads. The lowest reproduction accuracy images might be represented as "thumbnails" with images of increasing fidelity having increasingly larger reproduction sizes up to the highest reproduction fidelity at the original

image size. There are two reasons to do this. A page of small thumbnails can be rapidly scanned to decide which images to explore further. Secondly, small, low resolution images download quickly, allowing for rapid feedback in the navigation process. These two reasons are related, as they both involve navigability and bandwidth, both from network to screen, and for the user, from screen to decision. If a user can determine that they are not interested in an image from a glance at a thumbnail, both user and network resources have been saved.

Assuming the standard stochastic relaxation, the rate-distortion problem is then one of designing an algorithm or "code" for describing the data in a manner that will achieve the best possible tradeoff between the expected value of the rate (or per-symbol compressed description length) used to describe the data and the expected distortion achieved in the data reconstruction. Critical assumptions are that all expectations are taken with respect to the assumed underlying source distribution, and the distortion measure is assumed to be fixed and known at design time. While existing rate-distortion theory has many unresolved problems, recent advances by Effros et al. have begun the process of bringing together the theory and practice of *multiresolution* coding (e.g. [13]), which yield a single embedded description that can be read at a variety of rates. Even assuming near-optimal multiresolution codes, rate distortion theory fails to address any aspects of the geometry of Web layout, that is, why and where to put cuts and hyperlinks. The rest of this paper will explore new ways to approach this problem, and relate them to standard (lossless) source coding as formulated in Eq. (1).

III. ONE-DIMENSIONAL WEB LAYOUT MODELS

Users navigate through a Web site by browsing and clicking on, say, images at different resolution until they obtain the ultimate files they want. This is a complicated and highly individual process, and just as in standard source coding, almost any direct formulation of geometry layout will be intractable. Thus we will similarly seek an ensemble approach that captures the essential issues, emphasizing the large files that ultimately dominate the tail of the file transfer distributions. To that end we will assume that what is given is a collection of objects with their sizes and topological connection, and that any rate distortion coding has already been done, and focus particularly on multiresolution coded images. (We will ignore the details of this coding. Obviously, a research question of immediate interest is to do joint geometry and rate distortion coding.) We further assume that each object has some probability of access across an ensemble of users. This would naturally arise as users would tend to view, for example, a much larger number of thumbnails than high resolution images, and there might be non-uniformity in the probability of accessing different images at the same resolution. More generally, we would assume that the Web site *topology* is determined by the logical relationship between its component parts, which for a single progressively coded image file is obvious. What is not given by the content alone is the desired *geometry* of the Web layout, that is, essentially the specific locations of cuts in and hyperlinks between the images.

The assumed performance measure to be minimized through the Web site design is the average size of a downloaded file, exactly as in standard source coding. This is motivated by the

limitation on the bandwidth of both the network and the user, and particularly the average latency experienced in browsing through a Web site due to download time and decision time. In particular, it is highly desirable for the frequently accessed files that are primarily navigational to be small and download quickly (“mice”), as the users next action awaits this information, with large files (“elephants”) as the endpoint of the search process. The design degrees of freedom then are the cutting of progressively coded images into separate linked files, or generally the grouping of the objects into files. The result of design is the selection of the sizes of files and the locations of hyperlinks between them.

At one extreme, dividing the document into as many small files as possible lowers the cost of an individual download, but sacrifices the ease of manageability and navigability of the Web site. It also causes the Web server to process more requests for the same amount of information. Thus it is very important that the minimization of download time be subject to a constraint not only on the topology, but also on either the total number of files or on the total number, or average depth, or maximum depth of the hyperlinks. For most topologies, these latter constraints will be either exactly or roughly equivalent, and are motivated by the need for the Web site to be easily navigable by the user, with a reasonable number of clicks, and be maintainable by the Webmaster. This constraint is quite different from that of Kraft’s inequality in data compression, but plays a similar role in constraining the optimal solution. Thus in our model below, the number of files N is assumed to be fixed as a constraint from limited resources, like the constraint $\sum_i r_i \leq R$ in the PLR model.

A. Single Document Web Layout Design

Focusing on the simplest possible setting of a one-dimensional Web layout problem allows for analytical treatment. We begin with the case of building a Web site out of a single document such as a progressively coded image, then discuss the more complex scenario with multiple documents. Taking the stochastic ensemble approach, a real random variable X with sample space $[0, L]$ can be used to model user interest in various portions of the document, where L is the length of the document. The PDF $p(x)$ indicates the popularity of the information at location x . The specific functional form of $p(x)$ may vary with each individual document. Here we would make a reasonably generic assumption that the document is organized in such a way that more popular information is placed before less popular information, which means $p(x)$ is strictly decreasing with x , as shown in Fig. 1. This would certainly be the case for our assumed coded image.




Fig. 1. Document of size L .

Now suppose that we divide the document into N files, with the locations of the cuts being c_1, c_2, \dots, c_{N-1} . Denote the size of each file as l_i and the average popularity of each file as p_i . Then for any i , we have

$$\begin{cases} l_i = c_i - c_{i-1}, \\ p_i = \int_{c_{i-1}}^{c_i} p(x) dx, \end{cases} \quad (2)$$

where $c_0 = 0, c_N = L$. This division is illustrated in Fig. 2. Note that $\sum_{i=1}^N l_i = L$, and $\sum_{i=1}^N p_i = 1$.




Fig. 2. Dividing the document.

A one-dimensional Web site can be made by sequentially linking the above N files, so that each file corresponds to a Web page, as shown in Fig. 3. The design problem is to minimize the average size of files downloaded from Web site by optimally dividing the original document. To take into account the navigation behavior of users, we also assume that any user who visits the Web site has to start from the first page, and go through pages in ascending order. This means a user interested in page $\#i$ must access pages $\#1$ through $\#i$.

If p_i^t denotes the hit rate of file $\#i$ among all file transfers, then $p_i^t = \sum_{j=i}^N p_j$, because file $\#i$ is transferred whenever a user is interested in file $\#j$ with $j \geq i$. Note that p_i^t are not probabilities since $\sum_{i=1}^N p_i^t \neq 1$. However, we can say that for each file transfer, the download time is l_i with frequency p_i^t . Hence, minimizing the average downloaded file size means minimizing the cost function $J = \sum_{i=1}^N p_i^t l_i$ with respect to the locations of cuts c_1, \dots, c_{N-1} . To derive the optimality condition, let $\frac{\partial J}{\partial c_i} = 0, i = 1, \dots, N-1$. Then the following recursion relations hold:

$$p(c_i) l_{i+1} = p_i, \quad i = 1, \dots, N-1. \quad (3)$$

Since $p(x)$ is a strictly decreasing function, it is easy to verify that $l_i < l_{i+1}$ based on (2). This means the optimal file sizes will strictly increase with the index, but the exact values of the optimal l_i can only be solved numerically for general $p(x)$. If we suppose that users navigate through the document with a memoryless process, i.e., $P[X > x_0 + x | X > x_0] = P[X > x]$, then the exponential function is the only continuous PDF that meets this criterion. So assume that $p(x) = ke^{-\lambda x}$, where $0 \leq x \leq L$,




Fig. 3. A one-dimensional Web site.




Fig. 4. P_i^t vs. l_i for Web file transfers. ($p(x)$ is exponential.)

and $k = \frac{\lambda}{1 - e^{-\lambda L}}$. Let L be big enough so that $e^{-\lambda L} \ll 1$, and $k \doteq \lambda$. By substituting $p(x)$ into (3), we get the recursion equations for the optimal file lengths:

$$l_{i+1} = \frac{e^{\lambda l_i} - 1}{\lambda}, \quad i = 1, \dots, N-1. \quad (4)$$

The above relation indicates that the size of each file is exponentially larger than the one preceding it, which creates much faster than exponential growth in file size. Thus the last file in the sequence will be significantly larger not only than the first file, but also than the file immediately preceding it. The above equations can be solved numerically with the constraint $\sum_{i=1}^N l_i = L$. With the optimal solutions for l_i , we can compute $c_i = \sum_{j=1}^i l_j$, and furthermore, $p_i^t = \int_{c_{i-1}}^L p(x) dx = e^{-\lambda c_{i-1}} - e^{-\lambda L}$. Now normalize p_i^t such that $\sum_{i=1}^N p_i^t = 1$. If we consider the size of the transferred file as a random variable l , then the sample space is $\{l_1, l_2, \dots, l_N\}$, and p_i^t is the probability mass function (PMF) of l over each sample. The cumulative probability $P_i^t = P[l \geq l_i]$ vs. the size l_i is shown in Fig. 4, with $L = 10^4$, $N = 20$ and $\lambda = 0.1$.

The cumulative distribution of transferred file sizes displays an asymptotic behavior that is consistent with a heavy-tailed distribution. The slope of the curve in a log-log scale increases as the file size l_i grows bigger, and approaches -1 asymptotically as l_i becomes sufficiently large. This characteristic of the distribution is quite robust with respect to different parameter values N and λ as long as L is large enough. To prove this asymptotic behavior analytically, we need to consider the limit case where $L = \infty$. Then the probability of transferring file $\#i$ becomes $p_i^t = e^{-\lambda c_{i-1}}$. It follows that as $l_i \rightarrow \infty$,

$$\begin{aligned} \frac{\log p_{i+1}^t - \log p_i^t}{\log l_{i+1} - \log l_i} &= \frac{\log e^{-\lambda c_i} - \log e^{-\lambda c_{i-1}}}{\log \frac{e^{\lambda l_i} - 1}{\lambda} - \log l_i} \\ &= \frac{-\lambda l_i}{\log(e^{\lambda l_i} - 1) - \log \lambda l_i} \sim \frac{-\lambda l_i}{\lambda l_i - \log \lambda l_i} \sim -1, \end{aligned}$$

where $f(x) \sim g(x)$ denotes $\lim_{x \rightarrow \infty} \frac{f(x)}{g(x)} = 1$. The above argument shows that the slope of p_i^t vs. l_i on a log-log plot approaches -1 as l_i approaches infinity, which means $p_i^t \sim c l_i^{-1}$, $c > 0$. So the probability of transferred file sizes has a heavy tail with exponent $\alpha = 1$. Even for medium sized l_i , the

slope $\frac{-\lambda l_i}{\lambda l_i - \log \lambda l_i}$ is fairly close to -1 . And because $l_i \ll l_{i+1}$, $p_i^t \gg p_{i+1}^t$, the cumulative distribution P_i^t vs. l_i has exactly the same asymptotic behavior, which is verified by the numerical result demonstrated in Fig. 4.

To test robustness of the results to the assumed $p(x)$, we also considered one sided ($x > 0$) versions of both Gaussians and Cauchy distributions ($p(x) = \frac{\lambda/\pi}{1 + (\lambda x)^2}$). These two were chosen because they have very different types of tails, but are α -stable distributions. The Gaussian distribution has $\alpha = 2$ implying a finite variance, while the Cauchy distribution ($\alpha = 1$) is heavy tailed with infinite mean and variance. With these distributions, a simple recursion relation between l_i as the one in (4) is not available. However, similar asymptotic results, with similar proofs, on the tail of the probability of file transfers can be shown, and both asymptotics and numerical results are consistent with the exponential case. Thus power laws with $\alpha \approx 1$ are the robust outcome of this particular optimal design, independent of the assumed $p(x)$.

B. Aggregate Traffic from Multiple Web sites

The data with observed power laws ([6], [7]) uses aggregate web traffic involving a mixture of a large ($> 10^5$) number of files from different sites. The cumulative distribution plot of this data has a long tail with a similar slope to a line interpolating the final two points in Fig. 4, but the data would by comparison appear to have an almost continuous looking distribution, rather than such widely separated points. For a more direct comparison with the data, a modest extension to the single document case can be used to create a simple model of the distribution of a mixture of multiple web sites. This still allows for a largely analytic treatment, and can even produce a *continuous* distribution with power law tails that can be compared more directly with the data. Focusing again on images, suppose we consider the traffic from a large number of Web sites, where each has a homepage with thumbnails linked to a sequence of higher resolution images obtained from a progressive multiresolution coding. Neglecting the small homepages, which don't effect tail statistics, the result is a parallel sequence of chains, one for each image. Thus in this model there is no distinction between images on the same or different web site.

We'll assume that user interest within a coded image file, conditioned on selection of its thumbnail, is i.i.d. with an exponential distribution $p(x) = k_0 e^{-\lambda_0 x}$, as above. The recursion relation (4) can be used, so that each individual image file has a tail cut according to (4). What are the possible overall access distributions that are consistent with this assumption on the individual image documents? We could assume some distribution on the total sizes L and numerically compute the distribution of the resulting file sizes. However, if we simply assume we are given the distribution of the size and probability of access of the smallest file (l_1 above) in each chain, then we can very easily compute the distribution of all the file accesses in the Web site. Furthermore, if we make the abstraction that the l_1 are continuously distributed, we can compute a continuous distribution for the aggregate file transfer statistics.

For simplicity, assume that the documents are equally likely to be accessed, and that the distribution of l_1 is continuous with an exponential function on an interval I_1 with parameter λ_1 (its

PDF $f_1(l_1) = k_1 e^{-\lambda_1 l_1}$, $l_1 \in I_1$). This is an abuse of notation, with λ_1 determining the distribution of the individual l_1 (no longer a design variable, but now a given) while λ_0 determines the distribution of relative user interest in each document. We are further abusing notation and suppressing any explicit dependence on the different image documents. Let $h(x) = \frac{e^{\lambda_0 x} - 1}{\lambda_0}$. Then $l_i = h(l_{i-1}) = h^{i-1}(l_1)$, for all $i > 1$, where $h^m(\cdot)$ denotes the m th composite function of $h(\cdot)$. The inverse function of $h(x)$ is $g(x) = \frac{\log(\lambda_0 x + 1)}{\lambda_0}$, so $l_1 = g(l_2) = g^{i-1}(l_i)$. Then the PDF of l_i ($i > 1$) is

$$f_i(l_i) = f_1(l_1) \frac{\partial l_1}{\partial l_i} = f_1(g^{i-1}(l_i)) \frac{\partial g^{i-1}(l_i)}{\partial l_i}, \quad (5)$$

which is a complicated but computable function of l_i . The sizes of file transfers from the Web site can be viewed as a new random variable l_m , whose distribution is a mixture of all the distributions $f_i(l_i)$. And the mixing probability p_i^t is also a function of the particular value of l_i , which can be calculated as described in the previous subsection. Based on the theory of mixture probability, the cumulative distribution $P_m(x) = P[l_m > x]$ is given by

$$P_m(x) = \sum_{i=1}^N \int_x^\infty p_i^t(l_i) f_i(l_i) dl_i. \quad (6)$$

The above distribution is computed numerically for $N = 10$, $\lambda_0 = 0.25$, and $\lambda_1 = 1$ and the resulting P_m vs. l_m is displayed in Fig. 5 (dashed line). The upper tail of the distribution is well approximated by a power law with exponent ≈ 1.1 . To illustrate the cause of this heavy tail, the cumulative distributions of l_i conditioned on transferring file $\#i$ in the chain are also shown (solid lines). The plot is cut off at $x = 10^{10}$ to show the details of how the solid curves add up to give the dashed curve. Note that the continuous power law tail is created by the mixture of large files from different images, rather than from different files from the same image, as was the case in the discrete distribution in Fig. 4. The shape of the distribution is quite insensitive to the choice of N and λ_1 (and even the use of other distributions). As λ_0 increases the exponent of the power law tail approaches 1. This is an even more convincing correspondence between model and data than in Fig. 4, and it is encouraging that the model produces approximately the same $\alpha \approx 1$ as is seen in the data, and does so very robustly with respect to modeling assumptions.

C. Relevance to Web layout

What is particularly striking about this optimization formulation of Web layout is that it results not only in heavy tail access statistics, but robustly produces power laws with exponents $\alpha \approx 1$ close to what is observed in data. This is strongly preserved independent of the specific choice of the initial distribution of user interest $p(x)$ and when different Web sites are mixed into aggregate traffic. While we have for illustrative purposes focused on images, this framework should apply to other media and mixtures of media as well, although the analytic models may be more or less easily connected directly to the practical layout problem. Users navigating large text documents will typically browse a far larger number of reduced descriptions, such as titles and abstracts, than they will full documents. Video clips




Fig. 5. Cumulative distribution of l_m .

will have excerpts and still images, plus text descriptions for use in navigating to the ultimately desired large file downloads, and so on. Because the above models of both Web site structure and user behavior is highly simplified, we will also want to explore a variety of alternative models in this paper.

Note that this process of multimedia design has little to do with the Web per se, but should arise in almost any organization of information for human consumption that involves large objects. Thus, for example, it has been widely observed that libraries and file systems also have power law distributions of document and file sizes. Any process of accessing a large number of large papers, images, movies, books, or other media, necessarily requires a navigational process using reduced descriptions such as thumbnails, abstracts, reviews, excerpts, introductions, tables of contents, etc. Our analysis is relevant, provided the layout of such combinations of documents and navigational items can be described as the splitting or grouping of objects to minimize average file size, subject to a constraint such as limits on the clicks to reach large end destination files, or on the total number of files. It would be efficient for even a single text, video or audio stream to be cut into progressively downloaded pieces, if there was a reasonable chance that the user would stop after some intermediate point. The above analysis suggests that such cutting would naturally produce power laws with $\alpha \approx 1$, quite independent of the probabilities of the users stopping at any point in the file. A deeper insight, suggested by but not resolved by the analysis, is that the choice of media to create a successively more refined description would also lead to $\alpha \approx 1$, with text, images, audio, and video structured in a layered manner. Typically, one would expect text, audio, images, and finally video, to contribute at low to high file sizes, respectively, and this is apparent to some extent in the Web traffic data.

What is particularly new and different about the Web as a new media, however, is the existence of high quality data on access statistics, and the possibility to create a coherent theory. Of course, existing Web sites were not designed with this theory in mind, and individual Web sites are not likely and do not appear to have optimal layouts. Note, however, that the traffic statistics are for aggregate flows, and are thus mixtures of the statistics of different users accessing different Web sites, similar to the simplified mixture distribution considered above. For the aggregate flows to *appear* optimal, all that is required is that the variations across Web sites between optimal and actual be uncorrelated.

That is, traffic from individual Web sites need not have statistics consistent with our optimal layout theory, and they in fact do not appear to do so. They merely need to vary from optimal in such a way that these deviations average out when aggregate statistics are taken. Furthermore, Web sites that deviate substantially from this prescription would likely be so cumbersome and awkward to navigate that they either would be redesigned, or avoided, reducing their presence in the aggregate statistics. While there is no reason to expect individual optimality, there is equally no reason to expect systematic bias away from it, unless there is some very fundamental feature of Web layout not captured by our admittedly highly idealized abstraction.

One important insight to be gained from this research direction, even in its currently nascent state, is that the heavy tailed distributions characteristics of Web traffic (if not the detailed power laws) are likely to be an invariant of much of the future network traffic, regardless of the application. We expect that the current split of most traffic into elephants and mice will persist. (Note that the data from [2] was taken *before* the WWW was popular.) Most human-oriented communication processes that involve both active navigation and ultimately the transfer of large objects can naturally be “coded” this way. Even real-time, immersive virtual reality games or command and control systems and simulators are such that much of their traffic naturally codes into a combination of elephants containing configuration information and models together with mice that update the models in real time. Similarly, sensor and real-time control applications also naturally code into time-critical mice with measurement updates and actuator commands, against a background of elephants which update models of the dynamical environment and network characteristics. Of course, a coherent theory to make rigorous these informal observations is not yet available, and the HOT Web layout results are merely suggestive and encouraging. Nevertheless, we believe we have identified a theoretical basis for an important and empirically observed “invariant” of network traffic that must be taken seriously and treated systematically.

D. Generalized Source Coding

This subsection begins a deeper exploration of the origin of the specific $\alpha \approx 1$ in the power laws in the simplified models and in the data, and how they might change if alternative assumptions are made, or Web or other applications evolve. Some insight is gained from the approach in [10], which directly generalizes the source coding problem as formulated in (1) by using the loss function $l_i = f_\beta(r_i) = c(r_i^{-\beta} - 1)/\beta$, $\beta > 0$. Note that the standard case is recovered as $f(r) = -c \log(r)$ in the limit as $\beta \rightarrow 0$. When $f(r_i)$ has the more general form, the optimal solution to (1) is easily obtained analytically and leads to power law distributions with $\alpha \approx 1/\beta$. Here β can be interpreted an underlying “dimension.” Space limitations preclude a review of these results, but very briefly, the size of an object is associated with a β -dimensional volume, ξ^β , where ξ is a characteristic length scale of the object. The object size is limited by the resources, which can be thought of as $(\beta - 1)$ -dimensional cuts.

The Web layout considered above divides a one dimensional document into a chain of linked files corresponds to $\beta = 1$, and

standard data compression corresponds to $\beta = 0$, the only PLR case not producing power laws. Other design problems (such as firebreaks in forests, protective control in a power grid, rerouting flows in transport networks, etc) might correspond to other values of β . A comparison between the prediction from this PLR model with $\beta = 1$ and the 1995 data was also shown in [10] and the agreement is excellent. The 1998 data has $\alpha \approx 1.4$, which corresponds to $\beta \approx .7$. It was shown in [10] that $\beta < 1$ corresponds to a greater use of hyperlinks. This would be consistent with the change in the measured data from 1995 to 1998, which presumably corresponded to an increasing use of more sophisticated media and hyperlinks. Thus we could summarize our results so far as confirming the explanation in [10] that the Web traffic $\alpha \approx 1$ is due fundamentally to the fact that Web layout primarily involves cutting logically one dimensional objects (images, text, voice, video) into pieces for individual download. The one dimensional objects will typically not be some original image, text, video, or audio, but are the result of either some progressive multiresolution coding or the deliberate creation of reduced descriptions for navigational purposes. The role of additional hyperlinks between objects will be the focus of the remainder of the paper.

IV. SIMULATIONS OF GRAPH-BASED WEB LAYOUT MODELS

While the simple PLR model [10] yields surprisingly rich insights, the assumption that only the file size l_i depends on the resource r_i , while the probability of each file p_i remains fixed allows for an analytical solution but is a severe restriction, violated in Web layout design. It is easy to state the generalization of (1) to

$$J = \left\{ \sum p_i l_i \mid l_i = f(r_i), p_i = g(r_i), \sum r_i \leq R \right\} \quad (7)$$

but this fairly broad setting typically leads to a combinatoric optimization problem, which makes it extremely hard to solve analytically. Special cases may still allow analytic results, and the one-dimensional Web layout models introduced above is an example of this specialization. More complicated and realistic models of Web site design can be built on top of this, but there are no guarantees of finding the global optimum. Nevertheless, simulations of the optimization procedure using heuristic algorithms on graph-based Web layout models can be done and lead to results that are more realistic, yet still consistent with the simpler models.

In the next few subsections a series of graph models with increasing complexity and level of details will be described, including chains, trees, incrementally generated random graphs, and finally, a more sophisticated model with geometries in individual Web pages. While the tree and geometry models clearly involve more realistic structures, the random graph is introduced largely to explore the effect of other topologies, and further explore the robustness of this whole modeling approach. Once again, the results are entirely consistent with the simpler models above.

A. Web Sites with Chain Structure

To connect with the analytical results from the previous section, we start with the simplest possible graph, a chain-like

structure, where all the Web pages are sequentially connected by unidirectional links, as shown in Fig. 3. Let V_i represent each Web page and E_{ij} represent the hyperlink pointing from page V_i to page V_j . Assume that each V_i contains a single file with length l_i , which is a simplification from real Web pages since all the embedded objects (e.g., images) are considered as part of that file. We further assume that every user enters the Web site from the front page, V_1 , and at each page V_i has the choice of either following the hyperlink $E_{i,i+1}$ to browse the next page V_{i+1} with probability $p_{i,i+1}$ or exiting the Web site with probability $1 - p_{i,i+1}$. We will make the assumption that the exit probability p_e is fixed throughout, so that $p_{i,i+1} = 1 - p_e, \forall i$ and for all time. We make no distinction between a user returning to the homepage and restarting the search versus a new user starting at the homepage.

This navigation process can be modeled as a discrete-time Markov chain, where all the Web pages in the chain are considered as possible states for the random process, and the probability of going from V_i to V_j defines the transition probability p_{ij} . If there is no link connecting V_i and V_j , then $p_{ij} = 0$. It is easy to verify that this Markov chain is irreducible, aperiodic, and positive recurrent, therefore it is ergodic. For an ergodic Markov chain, each state V_i has a stationary probability p_i , which, in this model, corresponds to the probability of page V_i being accessed, or downloaded. So minimizing the average download time is equivalent to minimizing the cost function $J = \sum_i p_i l_i$, as in the model (7). Let $M = [p_{ij}]$ be the transition probability matrix. Then p_i can be computed by solving equation $p = M^T p$, where $p = [p_1, \dots, p_N]^T$ (column vector). For real Web sites, p_i could be directly calculated from log data with sufficiently large sample size. More advanced technology even allows the server to trace every link that a user follows, which provides information on the transition probabilities p_{ij} , or, the matrix M , from which p_i can be obtained.

Once the p_i have been computed, a simple heuristic is used to reduce the average file transfer J . The basic idea is to split high-hit files and merge unpopular ones. In particular, if a file is popular and large at the same time, split it into two smaller files. Repeatedly doing this will certainly lower J , but result in a Web site containing too many files to manage and to navigate through. Therefore, as a coarse approximation of the tradeoff between the cost and the manageability and navigability concerns, an upper limit is posed on the total number of files, as in the one-dimensional Web layout model. Alternatively, constraints could be put on the depth of the tree, the total number of hyperlinks, the outdegree, or any combination. For simplicity, we'll focus on constraining the number of files. To satisfy this, every time one file gets split up, two subsequent files that are rarely visited are chosen to be merged into a larger file. The splitting-merging procedure is repeated until the improvement of J is within a certain tolerance level or the number of iterations reaches a preset maximum.

That splitting popular files and merging unpopular would lead to heavy tails is not at all surprising. The question is how robust any resulting power laws are to variations in the process. For example, the criterion for choosing the page to split or the pages to merge can be different; while splitting, the location of the cut can either be in the middle, or picked at random; the param-

eters N and p_e can be chosen randomly in a certain range. We also need to initialize file sizes l_i to start the simulation, and three cases very far from power laws have been tested: the exponential distribution, the uniform distribution in $[l_{min}, l_{max}]$, and the degenerate case where $l_i = L$, for all i . Simulations on a number of chain models have shown that the simple optimization scheme works fairly well in continuously lowering the cost J . The improvement is quite dramatic at the beginning, then it slows down and finally reaches some steady state after some number of iterations. Let L^t be the random variable representing the sizes of transferred files. For every data set (l_i, p_i) , calculate the cumulative probability $P_i^t = \sum_{j \leq i} p_j$, with sizes of individual files ordered as $l_i \geq l_{i+1}$, i.e. $P_i^t = P[L^t \geq l_i]$. The resulting P_i^t from the optimized Web site always displays a power law decay with an exponent $\alpha \approx 1$. This result is consistent in all the simulations, and thus the one-dimensional chain robustly yields $\alpha \approx 1$, confirming our earlier results.

B. Web Sites with Tree Structure

A slightly more complex and realistic model for Web site layout than the chain is the rooted-tree, a simplification of the implicit hierarchical structure of many current Web sites. The first step of the simulation is to create the initial topology of the tree that reflects the intrinsic connectivity of the content. Two parameters are used to provide simple control over the key structural characteristics of a tree-like Web site: N_{node} , the number of nodes in the tree, and M_{out} , the maximum out-degree for all the nodes. Note that M_{out} and N_{node} together determine the minimum depth and the maximum degree of connectivity of the initial tree. The second step is to specify the Markov chain by assigning the initial transition probabilities p_{ij} . For each node V_i , ($i > 1$), let the exit probability be $p_{i1} = p_e$. The remaining probability $1 - p_e$ is evenly distributed among all the outgoing links E_{ij} ($j > 1, p_{ij} > 0$).

The splitting-merging algorithm for the tree-like Web sites is necessarily more involved than the one for the chains. In particular, the hyperlinks to and from the new Web pages need to be updated accordingly so that the tree structure is preserved throughout the optimization. Again small variations can be added to the basic algorithm in each simulation. Consistently, the cumulative distribution of file transfers from the optimized Web site follows a power law tail that varies from $\alpha \approx 1$ to $\alpha \approx 1.5$ depending on M_{out} . Larger M_{out} produces larger α , with the chain case and $\alpha \approx 1$ recovered in the case when $M_{out} = 1$. In general a tree with small M_{out} spreads out less and goes down more deeply, which is closer to the chain case with $\alpha \approx 1$. Thus larger α would be consistent with more use of nontrivial hyperlinks and tree structures, as opposed to pure linear chains.

C. Web Sites as Random Graphs

Both chains and trees can be viewed as special cases of more general graphs, and a random graph model is presented in this subsection, which will be referred to simply as the "graph model". Many aspects of the graph model are quite similar to those described in the chain and tree models, thus will not be repeated here. We will focus on the aspects that are unique to the random graph case. A random graph is not necessarily more realistic than a tree, but it serves to explore the impact of greater

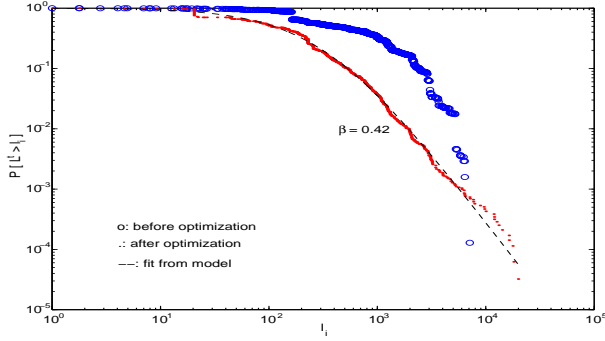


Fig. 6. P_i^t vs. l_i for Web file transfers before optimization and after 350 iterations, $N_{node} = 1000$, $N_{link} = 2$.

connectivity on the optimal distribution of file sizes and transfers.

Similar to the tree model, the graph model contains two structural parameters: N_{node} determines the scale of the Web site under study, and N_{link} , the number of incoming links at each page, governs the initial degree of connectivity between the Web pages. The graph model is built in an incremental way to crudely mimic the generation of real Web sites, where people initially put up some Web pages with hyperlinks in between, then gradually add more pages and connect them with the existing pages, and keep growing the Web site in an iterative fashion. Each node V_i has coordinates (x_i, y_i) , generated uniformly in a 10×10 square in \mathbf{R}^2 . The relevance of information between pages V_i and V_j is represented by their distance in the Euclidean space, the same idea used in [15]. The randomness in the locations of the nodes and consequently the connections between them are used to provide the degree of freedom in the topology. The resulting graphs typically have most nodes linked to nearby neighbors, with occasionally rather long links connecting nodes that are distant from each other.

The simulation result on a graph model with $N_{node} = 1000$ and $N_{link} = 2$ consistently produces tails with $\alpha \approx 2.3$, as seen in Fig. 6, which shows the cumulative distribution of file transfers P_i^t vs. l_i before and after the optimization. Although $\alpha > 2$ is not heavy-tailed, it does have an upper tail that is much heavier than the distribution before the optimization, which drops off exponentially at large file sizes. In fact, the optimization process pushes the whole distribution curve down at small and medium file sizes to reduce the average download time. As a tradeoff, the optimized Web site would produce more large file transfers than the original Web site.

The distribution of the sizes of *unique* files on the Web site, as opposed to the traffic from the Web site, also has a power law tail after the optimization. Let L^u be the random variable for the sizes of unique files, and P_i^u be its cumulative distribution, then $P_i^u = P[L^u \geq l_i]$. This yields $\alpha \approx 2$, a “heavier” tail than the distribution P_i^t for the transferred files. The heavy-tailed distribution for unique file sizes has also been observed from both the 1995 data and the 1998 data. In [6] it was shown that the set of file transfers is intermediate in characteristics between the set of file requests and the set of unique files. This is due to the use of caching in Web applications. At one extreme, if no caching is provided, then the set of transferred files should be the same

as the set of files that are requested, which is a simplification taken in this paper. On the other hand, in the case of perfect (infinite) caching, the set of files transferred should be identical to the set of unique files. In reality, caching does exist but is never infinite, so the distribution of resulting file transfers should be between the distributions of file requests and unique files. Back to the simulation results, P_i^t is less heavy-tailed than P_i^u because small files get requested more often than large files. If caching is taken into account, then the distribution of file transfers should have a $2 < \alpha < 2.3$ for this particular example. The effect of caching is not the focus of this paper, but will be an interesting subject for our future research.

These simulation results are quite robust with respect to the choices that can be made during the implementation of the optimization algorithm. Of course the objective function J may decay in different trajectories, and the resulting exponents of the power law tails can be slightly different, but the qualitative behavior of the graph model is fairly consistent in all the simulations. To capture the general characteristics of Web layouts with similar structure, the above design and optimization procedure was also simulated on a subnet consisting of K random Web sites, each with $N_{node} = 1000$ and $N_{link} = 2$. All the Web sites are considered equally probable to be accessed. So the overall statistics are just the average of those for individual ones. After the simulation all K data sets of (l_i, p_i) were mixed together and the only difference observed was that P_i^u have slightly smaller values of $\alpha \approx 1.8$. Additional simulations suggest that as the parameter N_{link} , the initial degree of connectivity of the Web site, increases, the distribution of Web file transfers has larger α , and ceases to be a power law at all.

The simulation results on chains, trees, and more random graphs with different N_{link} are all entirely consistent with the earlier analytic and numerical results, and suggests that the exponent of the power law tail in the distribution of file transfers from the optimized Web site is highly correlated with the degree of connectivity of the Web site. Among the three topologies that were studied, the chain structure has the least connectivity, the random graph with large N_{link} has the greatest connectivity, while the tree, similar to the random graph with $N_{link} = 1$, is between those two extremes. At the same time, the exponents of the tails show a similar pattern. Consistently, the more connected the graphs, the less heavy-tailed are the resulting distributions and the larger the α if there is a power law. This observation provides a possible explanation for why the power law tail from the 1998 data was steeper than that from the 1995 data. The reason could be that from 1995 to 1998 the design of Web layout became more mature and sophisticated, including more extensive use of hyperlinks that led to greater connectivity in resulting Web sites, which, in turn, would demonstrate less heavy-tailed distributions in file transfers.

D. Web Layout Model with Page Geometry

The Web layout models presented so far treat every Web page as an abstract node in a random graph without any geometry in it. Real Web pages are (mostly) HTML files containing embedded objects (images, etc.), distributed hyperlinks, and complicated layout. In this subsection a more realistic Web site model is introduced, which takes into account the internal structure of

each Web page. Since it is based on the random graph model, only the new features that are not in the previous model are described here. So each page V_i still contains a single file with length l_i , but each file is divided into one or more “paragraphs” or other objects to coarsely capture the logical connection between information. Different paragraphs are considered relatively independent, and it is ok to split them and move them into different pages when necessary. The number of paragraphs in each file and the locations of the paragraph breaks are generated randomly. In addition, the hyperlinks that point out of each page can be located anywhere in this page. To distinguish it from the graph model presented in the previous subsection, we refer to the model described here as the “geometric model”.

The optimization algorithm for the geometric model is more involved than for the graph model, because we need to keep track of the paragraph structure and the locations of the hyperlinks in the original Web pages. In particular, unless there is only one paragraph, the splitting point has to be between paragraphs. The relative locations of the hyperlinks in each paragraph will be preserved whenever files are split or merged. Space precludes a detailed discussion of our simulations with these models, but we can summarize our findings. A large number of simulations of the geometric model with different parameter settings have verified that the simple splitting-merging heuristic systematically reduces the cost function J . The distributions of transferred files (P_i^t) (and unique files (P_i^u)) from the simulations display power law tails, with exponents that vary from $\alpha \approx 1$ for linear chains, to $\alpha \approx 2$ for more highly connected random graphs. Again, this is consistent with all the earlier results and the data.

Although the geometric model has a higher level of realism than the graph model, and we do believe that these models capture some important features of Web layout design that are critical to the heavy tails in Web file transfers, there are many more characteristics of real Web sites that are still not included in our abstract models. The following are some of the issues that are worth looking into, in addition to ones already mentioned, such as combined rate distortion coding and layout, and constraints such as on depth, out degree, total number of links, in addition to simply the number of files. Obviously, the Markov chain navigation model is a simplification of perhaps more intricate user behavior. The validity of this approximation needs to be verified by empirical studies of user access patterns on real Web sites. In the case when user navigation demonstrates strong memory so that the Markov chain assumption cannot be used, the idea proposed in [16] that the average probability that a user visits the next page decays with time during each Web session can be adopted to modify our model. With this assumption the computation of the download probabilities p_i will involve repeatedly simulating user access to the Web site for a sufficiently large number of times.

There is no distinction between text (HTML) files and multimedia such as images, audio and video files in our models. Embedded objects are treated as part of a large file so that they can be reorganized during the splitting-merging procedure. However, multimedia files may sometimes be hard to split and merge, and for a complete treatment of this issue it seems essential that the rate distortion ideas be included. The paragraph structure in the geometric model captures this feature to certain degree. But

it is desirable to model embedded objects more explicitly so that their impact on the heavy tails of Web file transfers can be evaluated, especially since real-time multimedia traffic has become an increasingly larger contributor to the overall Web traffic. Note that according to a study by Crovella et al. in [6], although the presence of multimedia formats does add to the weight of the tail in the overall distribution of Web files, the distribution of text files itself is heavy-tailed. Finally, with the fast development in Web technologies today, many Web files are generated on real time when a user is browsing the Web site. It will be a challenging and interesting task for future research to study the impact of dynamic vs. static contents on Web site layout design and heavy-tailed Web traffic.

V. SUMMARY, DISCUSSION, AND FUTURE DIRECTIONS

This paper has presented a series of models of Web design which treated the layout of a Web site as an optimal design problem. The design objective was assumed to be the minimization of the average file size, and thus approximately the latency that the user experiences in downloading files while browsing the Web site, subject to constraints, such as on the total number of files. With either the analytic one-dimensional Web layout model or the simulation model for chain-like Web sites it was found consistently that the optimal distribution of Web file transfers had a power law tail with exponent $\alpha \approx 1$. Later we considered tree-like layouts and more random graphs and simulated the browsing and design process using heuristic optimization schemes. Again the distribution of file transfers from the optimized Web site displays a power law tail with exponent that increased with increasing hyperlinks, with a highly linked Web site having $\alpha \approx 2$. The observation that the distributions of Web file transfers become less heavy-tailed (larger α) with more complex Web layouts is also consistent with the change from the 1995 data to the data collected in 1998. An implication of this result is that optimal Web layouts and more effective use of hyperlinks may tend to produce slightly less bursty network traffic, but that the fundamentally one-dimensional nature of most media will tend to keep $\alpha < 2$.

Does the idealized problem described here actually explain the data, or is the remarkable agreement simply a coincidence? We have indicated tentative answers throughout this paper, but since this work is relatively new, we must caution that any answers are necessarily fairly speculative. It is possible to turn this question around however, and argue that even if the agreement with data is entirely coincidental, it is still true that Web sites *should* have these statistics if they are designed to minimize average file download times, and that this is an important design consideration. This has more serious implications and is in fact probably more important than the agreement with data. In particular, heavy-tailed distributions can be viewed as the *inevitable* outcome of a natural optimal “source coding” problem that is analogous to standard data compression, but with very different resulting distributions. Given the connection of heavy tails in Web sites and bursty network traffic, this can be thought of as bringing some initial closure to the origins of such traffic, but raises new questions. Bursty traffic is thus not an artifice of user behavior, but has some aspects which are intrinsic to at least the current dominant application, the WWW, and may be even

more intrinsic to any application which organizes information for human consumption.

The fact remains however that the correspondence between model and data is striking. We should, however, reemphasize (from Section III-C) that our approach's explanatory power does not depend on individual Web sites having "optimal" power law statistics, and they certainly do not. It only requires that the aggregate observed traffic, which is what matters for the network and servers, consist of individual sessions which mix to create the observed statistics. The "optimal" statistics are then a sophisticated null hypothesis about which individual Web sites and their access by users may vary because of reasons of aesthetics, content issues not captured in our models, poor design, or other idiosyncrasies. As long as there is no systematic bias in these variations, our approach explains the observed aggregate data. If there are systematic biases, then this model still represents a reasonable reference point to start explaining their origins. Ultimately any such biases should be included in more advanced models, but they should still produce the same power laws. Even the most complex Web layout models studied in this paper are still at a high level of abstraction from real Web sites. More realistic models of Web layouts with increasing levels of complexity can be built on top of these relatively simplified models. One next step is to gather empirical data of real Web site topology and user access patterns so that the relevance of the simulation models can be tested and enhanced. Informal exploration has so far been entirely consistent with our claims in this paper, but more systematic experimentation would be useful. Moreover, there are many more design objectives and resource constraints other than minimizing the average download time and limiting the total number of files on a Web site. Alternative settings for the optimization problem need to be explored to see whether similar conclusions can be drawn.

Even more important than the questions of data and model correspondence are what are the implications of these results for Web design, network traffic, network performance, and protocol design? If Web site layout can be viewed as "source coding" albeit with many strange and unfamiliar properties, then network protocol design and congestion control might profitably be viewed as a form of "channel coding", but also presumably with strange and unfamiliar properties. There may be some advantages in exploiting the specific features of the resulting source, as well as advantages in some joint design where the source coding reflects the nature of the network on which it must be transmitted. Recent work along these lines include designing new scheduling policies on Web servers for specifically dealing with heavy-tailed workloads [17]. The potential interaction between Web traffic modeling and caching is also interesting, since the use of Web caching affects the distributions of transferred files, while the design of more efficient caching strategies could exploit models of Web traffic. A discussion on this can be found in [7], as well as references to other papers on caching benefits and policies. There is also great potential in combining rate distortion techniques and Web layout design for more efficient transmission of information over the Internet. Another promising direction is to integrate generalized source coding with Low et al.'s optimization flow control theory [18] so that joint source and channel coding can be viewed as a global opti-

mization problem. Hopefully this mixed framework can offer a new look into different TCP congestion control algorithms under self-similar network traffic. We are investigating all of these issues in our current research.

Acknowledgement This work was supported by an AFOSR/DOD MURI for "Uncertainty management in complex systems," Caltech's Lee Center for Advanced Networking, and the EPRI/DOD Complex Interactive Networks program. Many of the authors cited herein also contributed directly to this research effort via numerous discussions. Special thanks to Michelle Effros, Deborah Estrin, Walter Willinger and the anonymous reviewers for insightful comments about the content of this paper.

REFERENCES

- [1] W.E. Leland, M.S. Taqqu, W. Willinger, and D.V. Wilson, "On the self-similar nature of Ethernet traffic," in *IEEE/ACM Transactions on Networking*, 1994, vol. 2(1), pp. 1-15.
- [2] V. Paxson and S. Floyd, "Wide-area traffic: the failure of poisson modeling," in *IEEE/ACM Transactions on Networking*, 1995, vol. 3(3), pp. 226-244.
- [3] W. Willinger and V. Paxson, "When mathematics meets the Internet," in *Notices of the AMS*, 1998, vol. 45(8), pp. 961-970.
- [4] A. Erramilli, O. Narayan, and W. Willinger, "Experimental queueing analysis with long-range dependent packet traffic," in *IEEE/ACM Transactions on Networking*, 1996, vol. 4(2), pp. 209-223.
- [5] W. Willinger, M.S. Taqqu, R. Sherman, and D.V. Wilson, "Self-similarity through high variability: statistical analysis of Ethernet LAN traffic at the source level," in *IEEE/ACM Transactions on Networking*, 1997, vol. 5(1), pp. 71-86.
- [6] M.E. Crovella and A. Bestavros, "Self-similarity in World Wide Web traffic: Evidence and possible causes," in *IEEE/ACM Transactions on Networking*, 1997, vol. 5(6), pp. 835-846.
- [7] P. Barford, A. Bestavros, A. Bradley, and M.E. Crovella, "Changes in web client access patterns: Characteristics and caching implications," in *World Wide Web: Special Issue on Characterization and Performance Evaluation*, 1999, vol. 12, pp. 15-18.
- [8] J.M. Carlson and J.C. Doyle, "Highly Optimized Tolerance: A mechanism for power laws in designed systems," in *Physics Review E*, 1999, vol. 60, pp. 1412-1428.
- [9] J.M. Carlson and J.C. Doyle, "Highly Optimized Tolerance: Robustness and design in complex systems," in *Physics Review Letters*, 2000, vol. 84(11), pp. 2529-2532.
- [10] J.C. Doyle and J.M. Carlson, "Power laws, Highly Optimized Tolerance and generalized source coding," in *Physics Review Letters*, 2000, vol. 84(24), pp. 5656-5659.
- [11] M. Newman, "Applied mathematics: The power of design," in *Nature*, 2000, vol. 405, pp. 412-413.
- [12] P. Bak, *How Nature Works: The Science of Self-Organized Criticality*, Copernicus, New York, 1996.
- [13] M. Effros, "Distortion-rate bounds for fixed- and variable-rate multiresolution source codes," in *IEEE Transactions on Information Theory*, 1999, vol. 45(6), pp. 1887-1910.
- [14] L. J. Schulman, "Coding for interactive communication," in *Special Issue on Codes and Complexity of the IEEE Transactions on Information Theory*, 1996, vol. 42(6), pp. 1745-1756.
- [15] B.M. Waxman, "Routing of multipoint connections," in *IEEE Journal on Selected Areas in Communications*, 1988, pp. 1617-1622.
- [16] B.A. Huberman, P.L.T. Pirolli, J.E. Pitkow, and R.M. Lukose, "Strong regularities in World Wide Web surfing," in *Science*, 1998, vol. 280(5360), pp. 95-97.
- [17] M. Harcol-Balter, M. Crovella, and S.-S. Park, "The case for SRPT scheduling in Web servers," in *MIT-LCS-TR-767*, 1998.
- [18] S.H. Low and D.E. Lapsley, "Optimization flow control I: Basic algorithm and convergence," in *ACM*, 1999, vol. 7(6), pp. 861-874.