Power laws, Highly Optimized Tolerance, and generalized source coding

John Doyle
Control and Dynamical Systems, California Institute of Technology, Pasadena, CA 91125

J.M. Carlson
Department of Physics, University of California, Santa Barbara, CA 93106

(March 16, 2000)

We introduce a family of robust design problems for complex systems in uncertain environments which are based on tradeoffs between resource allocations and losses. Optimized solutions yield the "robust, yet fragile" features of Highly Optimized Tolerance (HOT) and exhibit power law tails in the distributions of events for all but the special case of Shannon coding for data compression. In addition to data compression (DC), we construct specific solutions for world wide web traffic (WWW) and forest fires(FF), and obtain excellent agreement with measured data.

PACS: 05.40.+j, 64.60.Ht, 64.60.Lx, 87.22.As, 89.20.+a

In this Letter we introduce a class of optimization problems that we refer to as probability-loss-resource (PLR) problems, to reflect their three elements. PLR problems represent the simplest examples of *Highly Optimized Tolerance* (HOT), a mechanism for complexity based on robustness tradeoffs in systems subject to uncertain environments. The PLR problem is a generalization of Shannon source coding theory for data compression (DC), and solutions can be obtained analytically. As shown in Fig. 1, the resulting event size distributions agree quite well with experimental data for DC, as well as the well-studied, high quality data sets for World Wide Web traffic (WWW) [3], and forest fires (FF) [4]. These and other applications in which a single relationship between resources and losses applies over a broad range of scales in an otherwise relatively homogeneous substrate are particularly well-suited to the PLR setting. Further extensions of the PLR problem may explain the existence of heavy tailed distributions in a wide variety of complex systems in which design or evolution play a pivotal role.

The PLR objective is to minimize the expected cost

$$J = \{\sum p_i l_i \;\; | \;\; l_i = f(r_i), \quad \sum r_i \leq R\}. \tag{1}$$

Here i, $1 \leq i \leq N$, indexes a set of events, such as the occurrence of source symbols (DC), file accesses (WWW), and fire ignition and propagation (FF). Each event is assumed to be independent and initiated with probability p_i during some time interval of observation. Resources r_i are allocated to suppress the sizes l_i of the events, such that $l_i = f(r_i)$, subject to a bound R on the resources available. The l_i can be thought of as the loss or cost of an event, proportional to area burned (FF), or the lengths of files (WWW) and codewords (DC) to be transmitted on the Internet. A related continuous version with integrals replacing sums was considered in [1]. DC has a well-defined resource in codewords. In designing websites on the WWW, files and/or links plays a similar role. We assume the FF resources are firebreaks, roads, and other man-made mechanisms which limit fire propagation, and discuss natural occurring mechanisms.

We propose a general one-parameter resource vs. loss function $l_i = f_\beta(r_i)$ of the form

$$f_\beta(r_i) = \begin{cases} -c\log(r_i), & \beta = 0; \\ \frac{c}{\beta}({r_i}^{-\beta} - 1), & \beta > 0, \end{cases} \tag{2}$$

where resources are normalized so that $0 \leq r_i \leq 1$ and $f_\beta(1) = 0$, while the marginal loss per unit resource is ${f_\beta}'(r_i) = -c{r_i}^{-\beta-1}$, $\forall \beta \geq 0$. These conventions uniquely determine $f_\beta(r_i)$ to within the constant c. The strongest assumptions are that the p_i are independent of the resource r_i, and that the events are independent. Both are true in DC, but will be assumed throughout for simplicity.

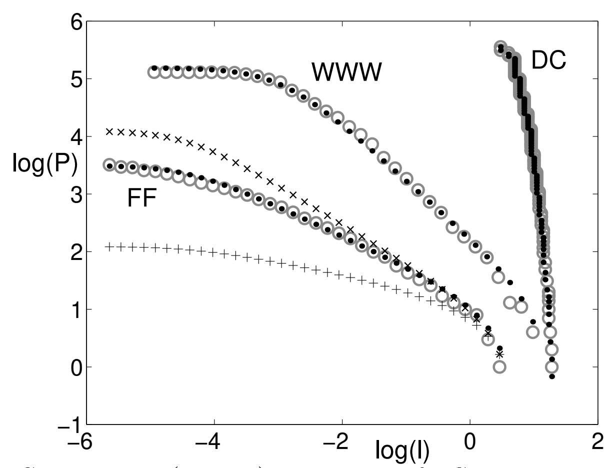

FIG. 1. Log-log (base 10) comparison of DC, WWW, and FF data (grey circles) with the results of the PLR problem (black dots) for $\beta = 0$, 1, and 2, respectively. Results for an SOC FF model with $\alpha = 1/\beta = .15$ (+) and an inaccurate PLR FF fit with $\beta = 3/2$ (x) are included for comparison. The cumulative distributions of frequencies $P(l \geq l_i)$ vs. l_i describe the areas burned in 4284 fires from 1986-1995 on all of the United States Fish and Wildlife Service Lands (FF), 130,000 web file transfers to 591 users on 37 machines at Boston University during 1994-95 (WWW), and codewords from data compression (DC). Both the size units (1000 km^2 (FF), 2 megabytes (WWW), and bytes (DC)) and the logarithmic decimation of the WWW and FF data are chosen purely for convenient visualization.

The values of β which characterize DC, WWW, and FF ($\beta = 0, 1, 2$, respectively) are discussed individually below. DC is a special case where $\beta = d = 0$ follows from a standard result in information theory, and the resource limitation is imposed by the need for unique decodability. Dimensional arguments motivate WWW and FF, where β is associated with the dimension d of the design problem. Suppose the loss or cost of an event is associated with a d-dimensional volume, ξ^d, where ξ is a characteristic length scale of the file accessed, or the region burned. The event size is limited by the resources, which can be thought of as $(d-1)$-dimensional cuts which isolate the event from the rest of the system. In WWW, dividing a one dimensional document into a chain of linked files corresponds to $d = 1$, while in FF dividing a two dimensional forest into areas corresponds to $d = 2$. Thus the resource density allocation in a given region is $r_i = \xi^{d-1}/l_i = \xi^{-1}$. This leads to $l_i \propto {r_i}^{-d}$, consistent with our interpretation that $\beta = d$ in the PLR problem (Eq. (2)). In general, β is determined by a resource/loss, which may or may not be directly related to physical dimension.

The optimal solution minimizes J (Eqs. (1–2)), and is obtained using standard constrained optimization methods (Lagrange multipliers). Setting the gradient of $\lambda(\sum r_i - R) + \sum p_i f_\beta(r_i)$ equal to zero yields $-p_i {f_\beta}'(r_i) = \lambda$, which equalizes the expected marginal loss and can be solved for the r_i. Then the optimal λ saturates the resource constraint with $\sum r_i = R$, $r_i \leq 1$ yielding

$$r_i = R p_i^{\frac{1}{1+\beta}} \left(\sum_j p_j^{\frac{1}{1+\beta}} \right)^{-1}, \tag{3}$$

$$l_i = \begin{cases} -\log(Rp_i) + \log(\sum_j p_j), & \beta = 0; \\ \frac{c}{\beta}\left[\left(R p_i^{\frac{1}{1+\beta}} \right)^{-\beta} \left(\sum_j p_j^{\frac{1}{1+\beta}} \right)^{\beta} - 1 \right], & \beta > 0, \end{cases} \tag{4}$$

$$J_\beta = \begin{cases} -\sum_i p_i \log(Rp_i) + (\sum_i p_i)\log(\sum_i p_i), & \beta = 0; \\ \beta^{-1}\left[R^{-\beta} \left(\sum_i p_i^{\frac{1}{1+\beta}} \right)^{1+\beta} - \sum_i p_i \right], & \beta > 0. \end{cases} \tag{5}$$

These globally optimized solutions provide a baseline for the study of applications such as WWW and FF, where resource/loss tradeoffs and design play some role, albeit less direct than is represented in the simple PLR setting.

Comparison with data: The data sets in Fig. 1 consist of pairs (l_i, P_i) with cumulative frequencies P_i of events of size $l \geq l_i$. In the standard "forward engineering" design problem the noncumulative p_i's are given, and optimizing the r_i produces the l_i in Eq. (4). "Reverse engineering" starts with the l_i data as given and generates model predictions using the PLR solution, with β given and c and R fit to the data.

Our aim is not to make detailed predictions, although these are surprisingly good, but rather to explore the qualitative differences in the data in Fig. 1. The DC data was generated by a standard DC algorithm that is equivalent to the PLR solution with the additional constraint of integer codeword length. Comparison with the WWW [3] and FF [4] data is less direct, as this data is aggregated from many individual websites and forests. The WWW and FF data start as large lists of event sizes, with presumably some measurement error in the FF case. The data in Fig. 1 has been logarithmically binned (decimated) for easier visualization, and ordered with $l_i > l_{i+1}$ and $P_i < P_{i+1}$, but our comparisons are insensitive to this binning.

A noncumulative distribution can be derived from the data using the difference approximation to $p = -dP/dl$

$$p_i = (P_{i+1} - P_i)/(l_i - l_{i+1}) \tag{6}$$

which can also be interpreted as the average probability density in the interval (l_{i+1}, l_i). This representation of the data can be compared with PLR predictions by inverting (4), which yields

$$\hat{p}_i(l_i) = \rho(\beta l_i/c + 1)^{-(1+1/\beta)} \tag{7}$$

for the probability density, where ρ is a constant set by the total number of events. The resulting comparison (not plotted in Fig. 1) between the model prediction and the data is good, although the data is noisy. Cleaner comparisons can be obtained using the cumulative distributions. In particular, Eqs. (6)-(7) yield

$$\begin{aligned} \hat{P}_i &= \sum_{j \leq i} \hat{p}_j (l_j - l_{j+1}) \\ &= \sum_{j \leq i} \rho(\beta l_i/c + 1)^{-(1+1/\beta)}(l_j - l_{j+1}), \end{aligned} \tag{8}$$

which we compute from the data $\{l_i\}$. The resulting PLR $(l_i, \hat{P}_i)$ pairs give strikingly accurate predictions when compared with the data (l_i, P_i) shown in Fig. 1. The shape of the body and tail of the curve is determined completely by β, with $\beta = 0, 1, 2$ for DC, WWW, FF, respectively, and is independent of ρ, c, and the details of the binning procedure. For $\beta > 0$, this leads to power law distributions ($P_i \propto l_i^{-\alpha}$) with $\alpha \approx 1/\beta$. The small scale cutoff depends on the constant c, which has units of loss l_i, and ρ shifts the curve vertically. The large scale cutoff is determined by l_1 in the data.

Both 2-d percolation and the standard self-organized criticality (SOC) Forest Fire model have $\alpha \approx .15$ [5] and this is plotted as well for comparison in Fig. 1 by setting $1/\beta = .15$. An additional "fit" to the FF data with $\beta = 3/2$ is also included to illustrate the substantial discrepancies associated with a mismatch in β. While the simple SOC FF model was not constructed specifically to produce the correct exponents, the alleged quality of the fits has nevertheless been used as an argument in support of the relevance of SOC models [4,7]. We considered the three additional FF data sets from [4], which yield similar excellent correspondence with the PLR predictions with $\beta = 2$, and discrepancies when compared with other $\alpha = 1/\beta$. We now consider the DC, WWW, and FF problems in more detail to derive the values of β used in these predictions.

Data Compression: Data compression has one of the simplest and most elegant design theories in all of engineering, and thus makes an excellent starting point. The objective in DC is to compress long source messages into short coded messages for more efficient storage or transmission [8]. The standard DC formulation due to Shannon [9] is the PLR problem with $\beta = 0$, although with different notation. 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 $\{0, 1, \ldots, D-1\}$. The resource limitation $\sum_i D^{-l_i} \leq 1$ (Kraft's inequality) is equivalent to the requirement that no codeword $c_i = y_1 y_2 \ldots y_{l_i}$ can be the prefix of any other codeword, which is necessary and sufficient for instantaneous decodability [8]. To see how the prefix condition leads to Kraft's inequality, let $x_i = 0.y_1 y_2 \ldots y_{l_i}$ be a real number in base D corresponding to c_i, and $[x_i, x_i + D^{-l_i})$ a subinterval of $[0, 1)$. The prefix condition implies all intervals are disjoint, since if any other codeword fell in that interval it would have c_i as a prefix, which is not allowed. Thus, $\sum_i D^{-l_i} \leq 1$. This argument can be reversed to construct an instantaneous code for any set of integer l_i satisfying $\sum_i D^{-l_i} \leq 1$.

Defining as the resource $r_i = D^{-l_i}$, the length of the ith interval, this exactly fits the PLR problem with $\beta = 0$ and $c = 1/\log(D)$. Minimizing the expected length of the coded message J (Eq. (1)) yields $r_i = p_i$ and $l_i = -c\log(p_i)$. The cost $J_0 = -\sum p_i \log(p_i)$ is the Shannon-Kolmogorov entropy. The data in Fig. 1 uses 16-bit source symbols and a standard scheme called Huffman coding to compress the postscript file of [2]. The slight discrepancy is due to integer codeword lengths, which we have neglected.

World Wide Web: Inspired by DC, we next cast efficient website design as a PLR problem. This requires more significant simplifying assumptions, in addition to approximating integer sizes (codewords or file lengths) by real variables. The simplest view is to treat website layout as the subdivision of a one dimensional document into files of length l_i, accessed with frequency p_i determined by user interest. The cost $J = \sum p_i l_i$ is roughly the average delay a user experiences in downloading files. The resource constraint preventing the use of many tiny files is that the hyperlinks between them consume space, user attention, and make web management and navigation difficult. This creates a tradeoff between simple websites and fast individual downloads, which we model as a resource constraint on the total number of links. In a one-dimensional chain of files this is equal to the number of cuts in the document.

The most subtle aspect of the PLR problem is determining β to reflect the resource vs. loss tradeoff. Subdividing a one dimensional document into linked files leads to an inverse relationship between the file size and the resource density, $l_i \propto {r_i}^{-1}$, so that $\beta = 1$ as discussed above. This does not account for the fact that in the division of the document the resource allocation r_i (separation of the document into linked files) may influence the p_i (hit probability for a given file). We can construct a microscopic model for which the PLR formalism (and the dimensional argument) applies exactly, by insuring that the p_i are not changed by varying the r_i. Formally, in $d-$dimensions we achieve this by splitting the system (in this case the one-dimensional document) into N regions of equal size L^d with uniform probability p_i of a user access in the ith region, and making the simplifying assumption that each such access is independent. We then design $(d-1)-$dimensional cuts subdividing the ith region into n_i equal regions (files) each of size l_i. Thus, the size of event i is $l_i = L^d/n_i$, while the resource allocation per unit loss is $r_i \propto (L/n_i)^{\frac{d-1}{d}}/(L/n_i)$. This yields $l_i \propto r_i^{-d}$, giving $\beta = d = 1$ for WWW.

The well-studied 1995 WWW data in Fig. 1 was measured when the web was in a nascent form, where files were primarily online versions of preexisting one-dimensional documents, consistent with $\beta = 1$. Subsequent web evolution to more efficiently exploit hyperlinks reduces the effective dimensionality, and this is consistent with more recent observations [11]. In our PLR setting, consider the extreme case in which every file in a "region" is linked to every other file, and links are still the limiting resource r_i. Then the number of links r_i scales as n_i^2, yielding $l_i \propto r_i^{-1/2}$ and $\beta = 1/2$. As in DC, more machinery is needed to develop useful tools for designing websites, and we have begun developing more complete models for website management. We have extended the PLR formulation to allow for the dependence $p_i(r_i)$, and with suitable choices very complex web topologies, resource limitations, and user navigation models can be incorporated in detail. The numerical solution of a variety of such problems with more realistic assumptions qualitatively match the results from the PLR problem with $1/2 < \beta < 1$ [10].

Forest fires: In FF we associate design with the subdivision of a two dimensional forest by one-dimensional man-made firebreaks and suppressors. This leads to $\beta = 2$, by the same $d-$dimensional argument constructed above for WWW. In FF the l_i are burned areas. The p_i are determined by the probability of sparks occurring in different regions and initiating fires. The cost J is the average timber lost in fires, which would tend to be minimized by deliberate design in managed forests. The resource r_i is the density of firebreaks and suppressors employed to stop the spread of fires. The tradeoff between use of land for trees or firebreaks sets the constraint on the total resources available.

The role of geometry and microscopic scales for FF are more complex and subtle than for WWW and DC, and we lack more sophisticated design methods to check the details. Thus our conclusions are necessarily much more speculative.

Interestingly, the exponents describing the distribution of recent fires do not seem to differ significantly from the historical record of fires in which there were no human interventions [4], so the PLR problem may be relevant to some extent to natural forests as well. The crucial fact may be that most mechanisms for fire spreading lead to expanding fronts. The fire terminates when the energy is absorbed by a resource, whether it be a firebreak (no fuel to burn), or some alternative man made or natural means of suppression. This much more general scenario still leads to $\beta = d$, with $d = 2$ for mesoscopically homogeneous forests. More effective fire prevention might shift the curve without altering the shape. Alternatively, landscapes which naturally break forests into regions of fractal dimension lower than 2 *would* have steeper power laws. For example, brush fires in California occur in unusually rugged terrain, and data [12] plotted similar to Fig. 1 is an equally good match to the PLR prediction with $\beta = 1$.

Discussion: While the correspondence between theory and data appears quite convincing, our theory is new, and clearly neither the 1995 web statistics nor the forest fire data arose from a purposeful global optimization of resource allocations such as we describe. Fig. 1 represents data from a large number of websites and forests over some limited time period, not a single source. A good fit is obtained as long as the deviations from optimality between different sites are roughly uncorrelated, a much weaker requirement than individual optimality. Furthermore, any individual websites with statistics far from optimal would simply present opportunities for practical application of the theory.

In percolation and other examples of equilibrium critical phenomena, as well as the standard SOC forest fire [5] and sand-pile [6] models, increasing the dimension d leads to steeper power laws, corresponding to a relative suppression of large collective fluctuations of microscopic degrees of freedom. In sharp contrast HOT consistently predicts exactly opposite trends. *Decreasing* the effective dimension in both the WWW and FF PLR formulations leads to steeper power laws, since for small β microscopic resources are more efficient in suppressing large events. Comparisons of the brush fire data [12] and the forest fire data [4], as well as the 1995 Web statistics [3] compared to more recent results [11] both lead to an effectively reduced dimensionality, and steeper power laws. All these examples support the statistical trends predicted by HOT, and deviate from those associated with criticality, independent of the specific values of the exponents. The HOT theory, and its application to DC, WWW, and FF, suggests that a new type of universality might apply to complex systems in which design and evolution play a role, but with most features in sharp contrast to familiar properties found in statistical physics [1,2].

As suggested by the name Highly Optimized Tolerance, 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 tradeoff leads to high performance and high throughput, ubiquitous power law distributions of event sizes, and potentially high sensitivities to design flaws and unanticipated or rare perturbations. HOT systems can be extremely sensitive (fragile) to design flaws associated with perturbations systems were not designed to handle (e.g. new viruses or invasive species) and where existing resources prove ineffective. In the PLR setting, inaccurate assumptions about the p_i for a known category of disturbance can result in misallocations of the r_i, sometimes with disastrous effects. This is particularly true for the small p_i, where few resources are allocated. Maximal costs associated with errors in p_i are of order the size of the largest event, so that the more tuned the design (i.e. the more nongeneric the allocation of r_i), the greater the performance, but also potentially the greater the risk.

ACKNOWLEDGMENTS

We thank Joshua Socolar for his insightful comments and suggestions. This work was supported by the David and Lucile Packard Foundation, NSF Grant DMR-9813752, a DOD MURI grant for "Uncertainty management in complex systems," Caltech's Lee Center for Advanced Networking, and EPRI/DOD through the program in Complex Interactive Networks.

[1] J. M. Carlson and J. Doyle, *Phys. Rev. E* **60**, 1412, (1999).
[2] J. M. Carlson and J. Doyle, *Phys. Rev. Lett.* **84**, 2529 (2000).
[3] M. E. Crovella and A. Bestavros, *IEEE/ACM Transactions on Networking* 5(6):835–846 (1997).
[4] B. D. Malamud, G. Morein, and D. L. Turcotte, *Science* **281**, 1840 (1998).
[5] K. Chen, P. Bak, and M. H. Jensen, *Phys. Lett. A* **149**, 207 (1990); B. Drossel, and F. Schwabl, *Phys. Rev. Lett.* **69**, 1629 (1992); K. Christensen, H. Flyvbjerg, Z. Olami, *Phys. Rev. Lett.* **71**, 2737 (1993).
[6] P. Bak, C. Tang, and K. Wiesenfeld, *Phys. Rev. Lett.* **59**, 381 (1987); L. Kadanoff, S. Nagel, L. Wu, and S. Zhou, *Phys. Rev. A* **39**, 6524 (1989).
[7] See, e.g., P. Bak, *How Nature Works: The Science of Self-Organized Criticality,* (Copernicus, New York, 1996).
[8] T.M. Cover and J.A. Thomas, *Elements of Information Theory*, (John Wiley and Sons, New York), 1991.
[9] C.E. Shannon, "A mathematical theory of communication," *Bell Sys. Tech. Journal* **27**, 379-423, 623-656 (1948).
[10] X. Zhu, J. Yu, and J. Doyle (unpublished).
[11] P. Barford, A. Bestavros, A. Bradley, and M. E. Crovella, in *World Wide Web*, Special Issue on Characterization and Performance Evaluation, Vol. 2, pp. 15-28, 1999.
[12] J. E. Keeley, C. J. Fotheringham, M. Morais, *Science* **284**: 1829-1832