Is Remote Host Availability Governed by a Universal Law?

John R. Douceur
Microsoft Research
johndo@microsoft.com

Abstract *– The availability of peer-to-peer and other distributed systems depends not only on the system architecture but also on the availability characteristics of the hosts participating in the system. This paper constructs a model of remote host availability, derived from measurement studies of four host populations. It argues that hosts are incompletely partitioned into two behavioral classes, one in which they are cycled on/off periodically and one in which they are nominally kept on constantly. Within a class, logarithmic availability generally follows a uniform distribution; however, the underlying reason for this is not readily apparent.*

1. Introduction

This paper presents a model that describes and partly explains the observed availability of remote computer hosts. For distributed systems in general, and for peer-to-peer systems in particular, a model of component availability is a prerequisite for understanding the availability of the overall system. Evidence from measured distributions of remote host availability strongly suggests that some common factors are at play. In particular, two phenomena seem to drive the dominant shape of the distributions: (1) an indistinct subset of hosts exhibit cyclical on/off behavior and (2) availability, particularly of non-cyclical hosts, tends to follow a uniform distribution.

For mathematical convenience, availability is stated in units of "nines," defined as $-\log_{10}$ of the fraction of time a host is not available. The term refers to the fact that, for example, a host that is available 0.99 of the time has an availability of $-\log_{10}(1 - 0.99) = 2$ nines.

Within this paper, the term *availability* is meant to encompass everything necessary for a host to be seen by another host: Its hardware must be powered on and operational; its software must be in a state from which it can respond to remote requests; and there must be an unbroken network route between the two hosts.

This paper examines availability distributions from four previously published studies of remotely observed host behavior, which are described in section 2. In these distributions, several common features are readily observable, and they are itemized in section 3. Section 4 shows that these features can be captured by a few simple graphical elements that can be fashioned into a mathematical model corresponding to a graduated mix of two uniform distributions. After section 5 shows that the model captures the dominant behavior of the observed distributions, section 6 offers and justifies the interpretation that the observed distribution arises because hosts have a probabilistic tendency to exhibit either cyclical or non-cyclical availability behavior. Section 7 muses on the as-yet enigmatic nature of the availability uniformity, and section 8 summarizes.

2. Measurements

Table 1 outlines the characteristics of four published studies of remote host availability: those of hosts at Microsoft taken in 1999 [1], of hosts on the Internet taken in 1995 [6], of hosts running Gnutella application software taken in 2001 [7], and of hosts running Napster application software taken in 2001 [7]. These measurements varied widely in their durations, polling frequencies, methodologies, and sampled populations.

Table 1: Details of Four Studies of Remote Host Availability

Data set	Microsoft	Internet	Gnutella	Napster
Date of measurement period	July 1999	1995	May 2001	May 2001
Duration of measurement period	5 weeks	3 months	60 hours	25 hours
Measurement polling interval	60 minutes	10 minutes	7 minutes	2 minutes
Measurement methodology	ICMP Echo	Sun RPC	TCP SYN	TCP SYN
Size of candidate host set	64,610	15,000	1,239,487	509,538
Criteria for inclusion in candidate set	on Microsoft corporate LAN	in random subset of DNS names	in dominant Gnutella overlay	storing popular song for Napster
Size of measured host set	51,662	1170	17,125	7000
Reason for reduction from candidate set	host attrition	host capability & admin tolerance	unspecified	unspecified
Data collected by	Bolosky et al.	Long et al.	Saroiu et al.	Saroiu et al.
Reference	[1]	[6]	[7]	[7]

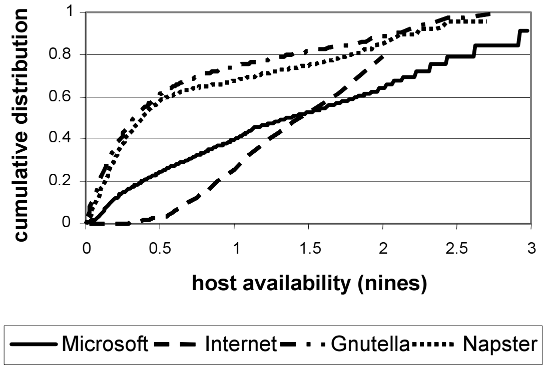

Fig. 1: CDFs of four host availability measurements

These studies presented their host-availability data in different ways. **Fig. 1** regularizes the presentations as cumulative distribution functions (CDFs), reporting availability logarithmically in units of nines.

None of the cumulative curves reach one, because there were hosts in all studies that never perceptibly failed during the measurement period.

3. Observations

Although there are significant differences between the four curves plotted in **Fig. 1**, there are also some noteworthy commonalities, highlighted in **Fig. 2**.

Observation 1: For availability values beyond about 0.7 nines, the dominant shape of the distribution is a straight line.

Observation 2: Extrapolating this linear region of the curve, the cumulative distribution reaches one in the availability range 2.4 – 3.4 nines.

Observation 3: For availability values less than about 0.2 nines, the dominant shape of the distribution may be linear, but it is hard to tell in this short range.

Observation 4: Within the availability range 0.2 – 0.7 nines, the distribution exhibits a pronounced bend.

It is also noteworthy that the Internet distribution's cumulative curve is concave up, but the others are concave down.

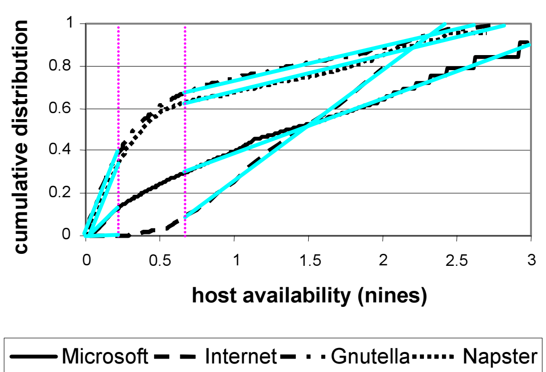

Fig. 2: Dominantly linear regions of curves in Fig. 1

4. Model

The simplicity of the common features among the CDFs suggests that they might be approximated by a simple analytical model. This section develops such a model, first by graphically mimicking the observed features of the measured distributions, then by formulating a mathematically precise expression for the graphical depiction, and finally by constructing a population that exhibits the model behavior.

4.1. Graphical presentation

Graphically, the proposed model's CDF includes two linear segments with a gradual transition between them, as illustrated in **Fig. 3a**. The low-availability segment is anchored at point (0, 0) and has slope s_l. The high-availability segment is anchored at a point (c, 1) and has slope s_h. If extended, the line segments intersect at an availability value of b, which is implicitly defined by s_l, s_h, and c.

The remaining characteristic of the model is the gradualness of the transition between the two linear segments, represented by the parameter g. This is more easily explained with reference to a probability density function (PDF), which is the first derivative of the cumulative distribution function, as shown in **Fig. 3b**.

The two straight lines on the CDF correspond on the PDF to horizontal lines with amplitudes equal to the slopes s_l and s_h. As the availability varies, the density shifts between s_l and s_h along a gradually varying curve centered at b. The gradualness of the bend in the CDF can be characterized by the maximum slope in the PDF, which occurs at point b. This maximum slope is proportional to g.

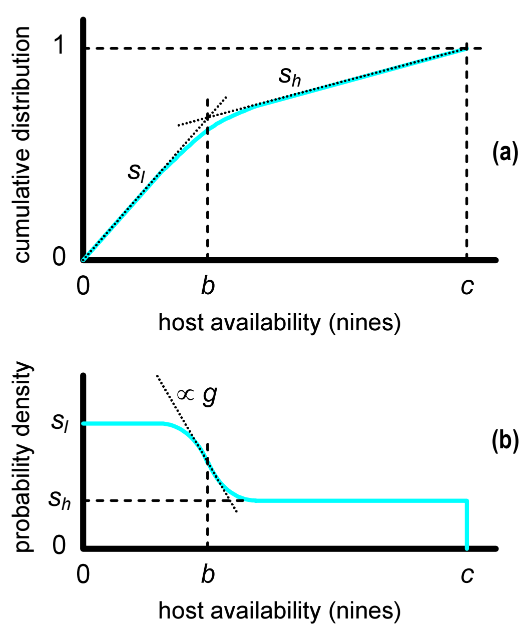

Fig. 3: (a) CDF and (b) PDF for analytical model

4.2. *Mathematical presentation*

If the PDF had displayed a step transition from s_l to s_h at point b, it could be mathematically expressed by employing a unit step function $u(x)$, which equals 0 for $x < 0$ and 1 for $x \geq 0$. The step-transition density function is thus

$$d_{\text{step}}(a;b,s_l,s_h,c) = s_l u(b-a) + s_h u(a-b) \quad (1)$$

for availabilities a satisfying $0 < a < c$ and 0 otherwise.

To capture the gradient in the PDF, the function $u(x)$ can be replaced by any symmetric, monotonic function that gradually shifts from 0 to 1 as its argument increases from $-\infty$ to ∞. An analytically simple function with these properties is the sigmoid function:

$$\sigma(x;g) = \frac{1}{1+e^{-gx}} \quad (2)$$

For this function, the gradualness is determined by g. Intuitively, g determines the maximum slope of the sigmoid curve, which occurs at $x = 0$:

$$\left.\frac{d\sigma}{dx}\right|_{x=0} = \frac{g}{4} \quad (3)$$

Replacing $u(x)$ with $\sigma(x;\, g)$, the expression for the density function becomes

$$d(a;b,g,s_l,s_h,c) = s_l\sigma(b-a;g) + s_h\sigma(a-b;g) \quad (4)$$

for $0 < a < c$ and 0 otherwise. As a density function, this expression is overspecified, because not all sets of parameters satisfy the constraint that the area under the curve equals one. A straightforward way to address this is to replace s_l and s_h with their ratio $r \equiv s_l / s_h$ and normalize by dividing out the total area between 0 and c. Expanding the expression for σ, the final expression for the density function becomes

$$d(a;b,g,r,c) = \frac{g\left(e^{ga}+e^{gb}r\right)\big/\left(e^{ga}+e^{gb}\right)}{gc+(r-1)\left(\ln\left(1+e^{gb}\right)-\ln\left(1+e^{g(b-c)}\right)\right)} \quad (5)$$

for $0 < a < c$ and 0 otherwise. Integrating this gives the cumulative distribution function

$$D(a;b,g,r,c) = \frac{ga+(r-1)\left(\ln\left(1+e^{gb}\right)-\ln\left(1+e^{g(b-a)}\right)\right)}{gc+(r-1)\left(\ln\left(1+e^{gb}\right)-\ln\left(1+e^{g(b-c)}\right)\right)} \quad (6)$$

for $0 < a < c$; 0 for $a \leq 0$; and 1 for $a \geq c$.

As an immediate corollary, if $r = 1$, this distribution collapses to a uniform distribution between 0 and c.

4.3. *Constructive presentation*

Understanding a complex distribution is often aided by viewing it as a combination of simpler distributions. For example, the Erlang distribution [3] results from a cascade of exponentials, and the Cauchy distribution [2] results from a ratio of two normals. Although a distribution can be decomposed in an infinite number of arbitrarily complex ways, this subsection presents a fairly simple way to view the distribution of Eqn. 5.

One way this distribution would arise is if:

- there are two classes of host availability behavior
- within each class, the distribution of logarithmic availability is uniform, and
- the probability that a host exhibits one class of behavior or the other is sigmoidally distributed

Stated concisely, the model distribution corresponds to a graduated mix of two uniform distributions.

Specifically, within one class, host availability is uniformly distributed in the range of 0 to c; within the other class, host availability is uniformly distributed in the range 0 to b. If each host definitively belonged to one of the two classes, the PDF would display a step transition at point b, as shown in **Fig. 4a**. Hosts in the first class would account for $s_h \cdot c$ of the probability mass, and those in the second class would account for $(s_l - s_h) \cdot b$ of the mass, respectively corresponding to the checked and hatched areas in **Fig. 4a**.

However, because hosts spend some fraction of their time in either class, the boundary is indistinct. In **Fig. 4b**, the checked area represents the contribution of hosts whose behavior is predominantly in the 0-to-c class, and the hatched area includes both hosts whose behavior is predominantly in the 0-to-b class and also those that exhibit no strong class ties. Grouping fickle hosts into the 0-to-b class is a somewhat arbitrary decision, but it facilitates the analysis in section 6.

5. Conformance

Fig. 5 (on the following page) shows the fit of the above model to each of the four measured availability distributions. Visually, the model adequately tracks the dominant shape of the measured curves. It does not pass Kolmogorov-Smirnov [4] goodness-of-fit tests, so if the model correctly accounts for the dominant host-availability behavior, then some other factor is responsible for the small but significant deviations.

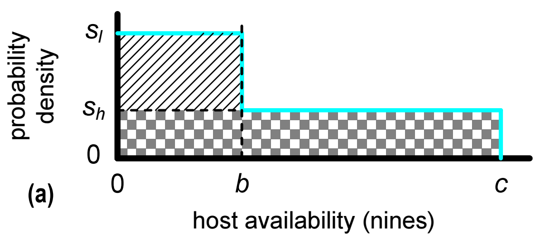

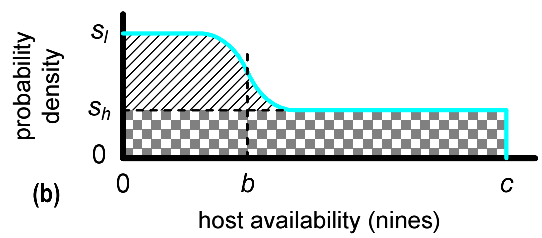

Fig. 4: Probability mass in PDFs for mixtures of two uniform distributions – (a) perfect mix (b) graduated mix

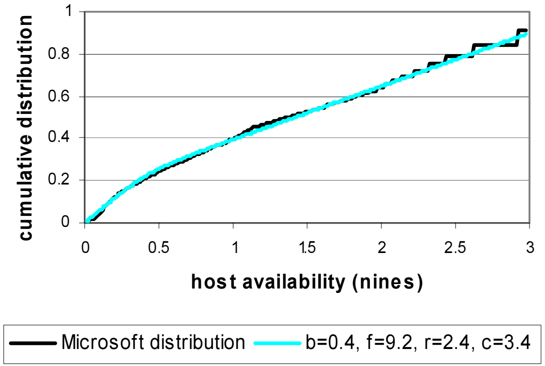

(a) 51,662 Microsoft hosts; 1999; 5 weeks; 60-minute intervals

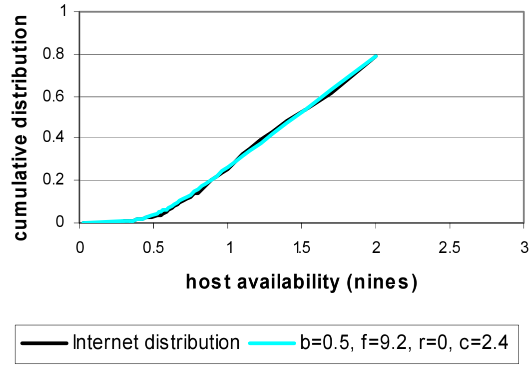

(b) 1170 Internet hosts; 1995; 3 months; 10-minute intervals

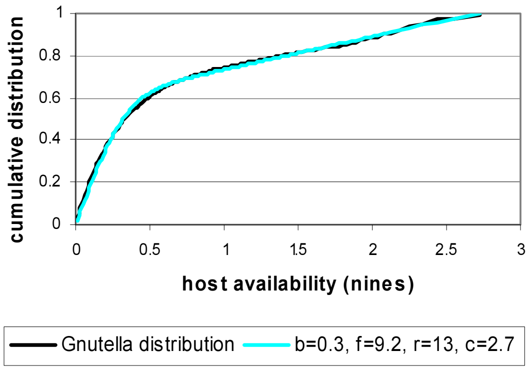

(c) 17,125 Gnutella hosts; 2001; 60 hours; 7-minute intervals

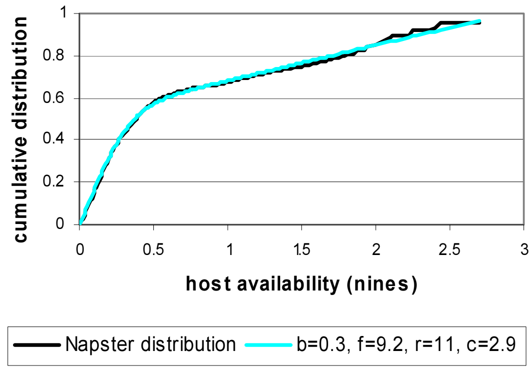

(d) 7000 Napster hosts; 2001; 25 hours; 2-minute intervals

Fig. 5: CDFs of four host-availability distributions and their analytical approximations

Table 2 shows the fitting parameters. It also shows the fraction of hosts that are not predominantly in the 0-to-c-availability set, calculated as:

$$f \approx (s_l - s_h)\, b = \frac{b\,(r-1)}{r\,b + c - b} \tag{7}$$

This fraction varies widely among the data sets and is the principal factor in shaping each distribution. More than half of the measured hosts that ran Gnutella and Napster software are outside the 0-to-c set, but only 14% of the measured hosts at Microsoft are.

For the Internet data set, f is negative, implying that the distribution is *not* a graduated mix of two uniforms. Rather, it is a graduated difference between uniforms: one uniform distribution from 0 to c, from which a sigmoidal region of low-availability hosts has been subtracted, as illustrated in **Fig. 6**. These hosts might have been removed from the population when Long et al. "…filtered to ensure that the hosts actually existed, could respond to the poll, and that their administrators would not mind the poll" [6]. If these criteria were positively correlated with host availability (which the verifiability of host existence clearly is), such filtering could account for the removal of low-availability hosts.

Table 2: Fitted model parameters for four data sets

Data set	b	g	r	c	f
Microsoft	0.4	9.2	2.4	3.4	0.14
Internet	0.5	9.2	0	2.4	–0.26
Gnutella	0.3	9.2	13	2.7	0.57
Napster	0.3	9.2	11	2.9	0.51

Perhaps the most striking result is that the fitted value of g is the same for all four distributions. This parameter relates to the curvature of the CDF insofar as the maximum slope of the PDF is $(s_h - s_l) \cdot g/4$. In words, the sharpness of the transition between the two flat regions of the PDF is proportional to the difference between the amplitudes of the regions multiplied by a seemingly arbitrary constant. However, this constant is only arbitrary in the sense that it is determined by the logarithmic base for calculating availability. If the natural log base e had been used instead of the decimal base 10, g would equal 9.2 / ln(10) = 4, so the constant in the expression for the PDF's maximum slope would be one. Thus, if the model behind Eqns. 5 and 6 were more parsimonious, g would be extraneous.

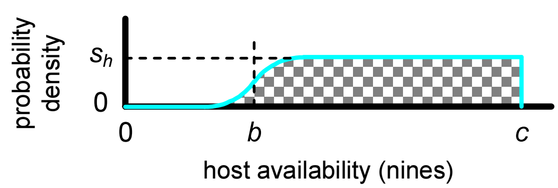

Fig. 6: Uniform PDF with low-availability hosts removed

6. Interpretation

Why should there be two classes of hosts with different availability characteristics? One hypothesis is that hosts in the 0-to-b class are periodically cycled on and off, and hosts in the 0-to-c class are nominally left on at all times. This seems plausible because, in the measured distributions, the boundary value b is in the range of 0.3 to 0.5 nines, indicating hosts that are generally on for less than half ($1–10^{–0.3} \approx 1/2$) to two-thirds ($1–10^{–0.5} \approx 2/3$) of the time, which is consistent with being turned off overnight or over weekends.

If this hypothesis is correct, then **Table 2** implies that the Internet data set should display no cyclicality, a conjecture that could be verified given the raw data. In theory, the Gnutella and Napster data sets should show a large degree of cyclicality; however, this might be difficult to observe within the short time span over which these measurements were made.

According to **Table 2**, the Microsoft data set should show significant cyclicality in 0.14 of the measured hosts. This can be validated as follows: For each host, compute a Fourier transform of the uptime sequence and calculate the fraction of energy in the daily or weekly spectral bands. Sort the hosts by this cyclical energy fraction; remove the top t of the hosts from the data set; and plot the availability distribution of the resulting set, as in **Fig. 7**. The distribution for t = 0.14 is very close to the distribution $D(a; \cdot, \cdot, 1, 3.4)$, which it should be if the model and hypothesis are correct.

7. Conundrum

The tendency of hosts to exhibit either cyclical or non-cyclical behavior seems intuitively reasonable, but it is not at all clear why logarithmic availability should be uniform. Although none of the measurement studies ran for long enough to observe the behavior of the most available hosts, three of the cumulative distributions come rather close to one, and they exhibit no obvious deviation from uniformity. Irrespective of whether this trend continues all the way to a cumulative frequency of one, the general phenomenon wants explanation.

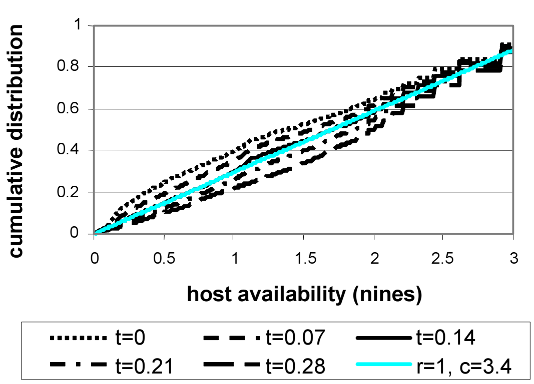

Fig. 7: Microsoft data sans *t* most cyclical hosts

8. Summary

This paper examined availability distributions from previously published studies of remotely observed host behavior, noted several common features among the distributions, and proposed a graphical, mathematical, and behavioral model to account for the dominant shape of the distributions. The model posits that hosts probabilistically exhibit availability behavior from each of two classes, one in which they are nominally left on at all times and one in which they are periodically cycled on and off. Within each class, host availability follows a uniform distribution. In one of the published distributions, aggressive population filtering appears to have eradicated nearly all low-availability hosts, yet the model still matches the distribution closely.

The potential value of this model is in providing an approximation of and a partial explanation for observed host availability behavior, akin to laws observed by Lotka [5] and Zipf [8]. However, the behavior underlying the uniform regions of the distribution remains unexplained.

Acknowledgements

The author is indebted to Jon Howell for his comments on numerous incomplete and poorly written versions of this paper as its core ideas took shape, to Jay Lorch for his indefatigable skepticism that prompted a radical restructuring of the paper into its current form, and to Stefan Saroiu for providing access to his Gnutella and Napster host uptime distributions.

References

[1] W. J. Bolosky, J. R. Douceur, D. Ely, M. Theimer, "Feasibility of a Serverless Distributed File System Deployed on an Existing Set of Desktop PCs", *SIGMETRICS 2000*, 2000, pp. 34-43.

[2] J. E Freund, *Mathematical Statistics, 5th ed.*, Prentice Hall, 1992.

[3] L. Kleinrock, *Queueing Systems, Volume I: Theory*, John Wiley & Sons, 1975.

[4] D. E. Knuth, *The Art of Computer Programming: Volume 2, Seminumerical Algorithms, Second Edition*, Addison-Wesley, 1981.

[5] A. J. Lotka, "The Frequency Distribution of Scientific Productivity", *Journal of Washington Academy of Sciences* 16(12), 1926.

[6] D. D. E. Long, A. Muir, R. A. Golding, "A Longitudinal Survey of Internet Host Reliability," *14th SRDS*, IEEE, 1995, pp. 2-9.

[7] S. Saroiu, P. K. Gummadi, S. D. Gribble, "A Measurement Study of Peer-to-Peer File Sharing Systems," *MMCN '02*, SPIE, 2002.

[8] G. K. Zipf, *Human Behavior and the Principle of Least Effort*, Addison-Wesley, 1949.