LETTER

On Temporal Locality in IP Adress Sequences

Weiguang SHI[†a)], *Nonmember*, **Mike H. MACGREGOR**[†], *Member*, *and* **Pawel GBURZYNSKI**[†], *Nonmember*

SUMMARY Temporal locality in IP destination address sequences can be captured by the addresses' *reuse distance* distribution. Based on measurements from data for a wide range of networks, we propose an accurate empirical model in contrast to results derived from the stationarity assumption of address generation processes.
key words: *temporal locality, LRU stack, reuse distance*

1. Introduction

An *LRU (Least Recently Used) stack* [1] is a vector of distinct data items, IP addresses in our case. Each address in a sequence is looked up in the stack. If the address is not found it is placed on top of the stack and the addresses in the stack are shifted down by 1 position. If the address is found, it is put on the top of the stack and addresses that were previously above it are shifted down by 1 position. The index at which the address is found (0 refers to the stack top), in the stack, is called the stack distance, or *reuse distance*. In the case that the address is not in the stack, its reuse distance is defined as the stack size. Because a smaller reuse distance indicates the current address was referenced a short time ago, the reuse distance distribution quantifies the concept of temporal locality, which refers to the phenomenon that when a data item is referenced, it is to be referenced in the near future.

Ref. [2] derives the reuse distance probability distribution by assuming the stationarity of the underlying address generation process. The inverse stack growth function (ISFG), $f(t,\tau)$, equivalent to the average working set size function [3], relates the size of the stack (the number of distinct addresses) to discrete time t and the working set window size τ. In addition, the stack growth function (SGF), g, is defined as f^{-1}. The stationarity assumption leads to:

$$f(t1,\tau) = f(t2,\tau) \tag{1}$$

where $t1$ and $t2$ are any two instants in time. Thus, the stack size is dependent only on τ. $f(\tau)$ is found to follow a power law:

$$f(\tau) = \tau^{\alpha} \qquad (\tau \gg 1). \tag{2}$$

Manuscript received March , 2003.
[†]The authors are with Computing Science Department, University of Alberta, Canada.
a) E-mail: wgshi@cs.ualberta.ca

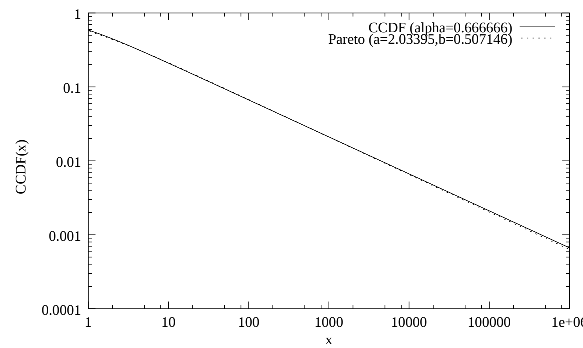

Fig. 1 The CCDF in Eq. 4 matches the Pareto in Eq. 5

The probability of the reuse distance k, a_k, is proved to be

$$a_k = \{f(g(k-1)+1)-(k-1)\} \\ -\{f(g(k)+1)-k\}. \tag{3}$$

2. The CCDF of Reuse Distances

Eq. 2 describes an asymptotic behavior that may not apply to small τ; this could lead to inaccurate predictions for the small reuse distances that appear frequently as the result of temporal correlation in IP address sequences. For this reason, Eq. 3 is not acceptable for the performance evaluation of finite-size caches. We will show that the reuse distance distribution predicted by Eq. 3 differs significantly from empirical measurements.

The complementary cumulative distribution function (CCDF) for reuse distances predicted by the ISFG can be derived from Eq. 3:

$$CCDF(k) = 1 - \sum_{i=1}^{k} a_i \\ = (k^{1/\alpha}+1)^{\alpha} - k \tag{4}$$

This CCDF is a straight line in a log-log plot (Fig. 1). For $\alpha = 2/3$ [2], the CCDF curve of the Pareto distribution:

$$P(x) = (ax+1)^{-b}, \quad a,b > 0 \tag{5}$$

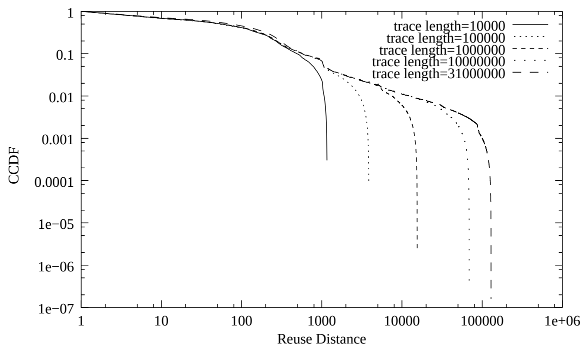

Fig. 2 Measured reuse distance CCDF's

with $a = 2.03395$ and $b = 0.507146$, nearly matches the curve of Eq. 4. From Fig. 1, one may conjecture that any reuse distance CCDF in the form of Eq. 4 can be approximated by some Pareto CCDF. For large x, by binomial expansion for real exponents, Eq. 4 and Eq. 5 lead to:

$$CCDF(x) = \alpha x^{1-1/\alpha} + \alpha(\alpha-1)/2x^{1-2/\alpha} + O(x^{1-3/\alpha}) \quad (6)$$

and

$$P(x) = (ax)^{-b} - b(ax)^{-b-1} + O(x^{-b-2}), \quad (7)$$

respectively. If $b = 1/\alpha - 1$ and $a = \alpha^{-1/b} = \alpha^{\alpha/(\alpha-1)}$, then $P(x)$ is a good approximation to $CCDF(x)$ when x is large.

3. Empirical Measurement Results

Fig. 2 shows measured reuse distance CCDF's for sub-traces extracted from a 31-million-entry trace, LDestIP, published by National Laboratory of Applied Network Research (NLANR)[4]. Each curve represents a CCDF calculated from a sub-trace of a certain length. All sub-traces are longer than 10,000 entries. It is evident that each curve comprises of several distinct segments and can not be adequately fitted by a straight line in a log-log plot.

The CCDF's follow a pattern. Initially, the curves are all very close. This consistency across different lengths of the traces implies stationarity. Over 60 percent of reuse distances fall in this range. Afterwards, the curves for different lengths diverge. We call this part of the curve the tail. The tails start with segments of roughly straight lines, though there are bumps in the middle of these segments. Then they drop off nearly vertically. The diverging tails give the look of branches from a common trunk.

Measurements show that the final segments, i.e., the almost vertical parts of the CCDF tails in Fig. 2 are linear. This linearity is due to addresses that are not *reused*, but appear only once in the traces. Using a longer trace to calculate the CCDF eliminates the linear tail of the CCDF for a shorter trace, although this introduces another longer linear tail. The reason is that the addresses that appear only once in a short trace can appear more than once in a long trace. The phenomenon that linear branches stem from a common base in all the CCDF figures reflects instability due to the limited length of the traces. In other words, if we had infinitely-long traces, there would be no linear tails in the CCDF's and they would converge to one common shape.

4. The Proposed Model

After removing the nearly vertical segments of the curves, which are caused by addresses that appear only once in the traces and thus are due to the limited trace length, we have been able to fit the CCDF's with the mixed distribution (Eq. 8):

$$C(x) = pW(x) + (1-p)P(x),\ 0 < p < 1, \quad (8)$$

where $W(x)$ is the CCDF of the Weibull distribution

$$W(x) = e^{-(x/d)^c}, \quad c, d > 0,$$

and $P(x)$ is the CCDF of the Pareto distribution as shown in Eq. 5.

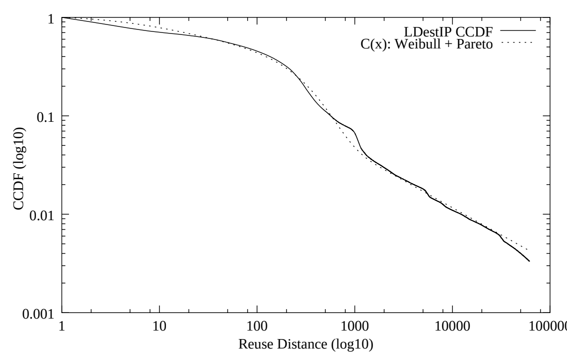

Fig. 3 Fitting a reuse distance CCDF with Eq. 8

Fig. 3 shows the fit for the LDestIP trace. We have experimented with a large number of traces for networks ranging from campus level networks to major Internet backbones, and have found that Eq. 8, with appropriate parameters, is able to closely match the observed reuse distance CCDF's.

To achieve parsimony, we have tried to replace the Weibull component in Eq. 8 with the CCDF of an exponential distribution. The results are promising and this reduces the number of parameters needed to four. Furthermore, the discussion in the previous sections indicates that we might be able to replace the Pareto component in Eq. 8 with a function in the form of Eq. 4, possibly reducing the total number of parameters to three. Experiments are under way.

5. Conclusion

In this Letter, we describe the discrepancies between theoretical results and empirical observations of reuse distance distribution, a measure of temporal locality in IP destination address sequences; we propose an accurate and flexible mixed distribution model to characterize real world locality.

Acknowledgment

The authors wish to thank Professor Robert Israel at University of British Columbia for helping to describe the relationship between the Pareto CCDF and Eq. 4.

References

[1] R. Mattson, J. Gecsei, D. Slutz, I. Traiger. Evaluation techniques for storage hierarchies. *IBM Systems Journal*, 9(2):78–117, 1970.
[2] M. Aida, T. Abe. Stochastic model of Internet access patterns. *IEICE Transactions on Communications*, E84-B(8):2142–2150, 2001.
[3] P. J. Denning, S. Schwartz. Properties of the working set model. *Communications of the ACM*, 15(3):191–198, 1972.
[4] NLANR (National Laboratory for Applied Network Research) Measurement and Operations Analysis Team (MOAT). Packet header traces. http://moat.nlanr.net.