University of Alberta

Locality in Internetwork Traffic

by

Michael H. MacGregor
and
Ivan L. Chvets

Technical Report TR 02-04
March 2002

DEPARTMENT OF COMPUTING SCIENCE
University of Alberta
Edmonton, Alberta, Canada

Abstract

As the amount of traffic carried by the Internet grows, the rate at which packets must be processed at each router and switch in the network also increases. The results of routing table lookups can be cached in order to accommodate this increasing demand for fast delivery of packets. However, there must be some locality in the traffic in order for caching to be efficient. This paper presents two methods for analyzing temporal and spatial locality in IP traffic. A high degree of locality is found in both dimensions. It was also found that the usual mapping of IP addresses used to preserve confidentiality destroys the address distribution and spatial locality of the trace.

1 Introduction

As the amount of traffic carried by the Internet grows, the rate at which packets must be processed at each router and switch in the network also increases. The results of routing table lookups can be cached in order to accommodate this increasing demand for fast delivery of packets. However, there must be some locality in the traffic in order for caching to be efficient. The goal of the work described in this paper is to investigate the types of locality in the traffic carried by large internetworks, and to quantify the amount or degree of such locality.

This paper presents the results of several analyses of locality in IP traffic. First, related work in the general area of network address locality is discussed. Then several new algorithms for analyzing locality are presented along with results validating the algorithms. Finally an analysis of some traffic traces is presented along with the results produced by the algorithms.

2 Related Work

Some previous research has been conducted in the area of locality of reference in network traffic. However, the type(s) and degree(s) of locality in large IP-based internetworks has not been characterized.

[Jain 90] documents the characteristics of destination address locality in local area networks as well as a comparison of different caching techniques. This study showed that the pattern of references to destination addresses in local area networks has high temporal and spatial locality. However, traffic patterns in local area networks are very different from those in large internetworks. Thus, this study cannot be used as proof of the existence or degree of locality of address references in IP traffic in large networks.

Another study [Claffy 94] identified destination address locality as one of the important characteristics of internetwork traffic and showed a case where traffic flow was very non-uniform: over 50% of the load was destined to only 2.8% of the network addresses. However, this study did not identify the types of locality which are present in the traffic, nor did it suggest how destination address locality could be exploited.

The study documented in this paper was an investigation of traffic locality in large IP-based internetworks. Methods for exploiting locality to improve the performance of routers and switches are suggested.

3 Overview of Locality

The research conducted in the area of memory and file referencing behaviour in computer systems has shown that these reference patterns follow non-uniform distribution in time and space. In other words, some memory locations or parts of a file are being accessed more frequently than others. To describe this behaviour term "locality of reference" has been introduced. The locality of reference has been exploited in various ways to improve performance of computer systems. One technique is to use a small fast memory (cache) to store the most frequent accesses in order to speed up information retrieval. This section presents an overview of the basic concepts of locality of reference.

Locality of reference has been studied in the context of memory reference behaviour [Bunt et al 84], file systems and virtual memories [Denning 70]. Generally, the principle of locality of reference can be described as follows [Hennessy et al 90]:

- If an item is referenced, it will tend to be referenced again in short period of time (*temporal locality*)
- If an item is referenced, neighbouring items will tend to be referenced soon (*spatial locality*)

There are identified two types of locality: *temporal locality* and *spatial locality*. Temporal locality suggests that the information last referenced has a high probability of being referenced again in the near future. For example, references performed by a program in a loop have a high degree of temporal locality, because data and instructions in the loop are being re-used. Spatial locality implies a high probability of referencing *neighbouring* regions of the region last referenced. For example, accessing elements of an array sequentially would produce a string of references with a high degree of spatial locality, because elements of the

array are stored in memory at adjacent or neighbouring locations. The notion of neighbouring addresses for memory and file system addressing schemes is clear; however, for computer networks it requires some additional clarification. Addresses in computer networks can be considered in the same neighbourhood if they belong to some address group created by the implementation of some predefined addressing scheme for division of the network address space. These address groups are called subnets.

The terms *persistence* and *concentration* have also been used to characterize locality behaviour [Bunt et al 84]. Persistence refers to the tendency to repeat references to a single address. This is related to temporal locality. Concentration suggests a tendency for references to be limited to a small group of addresses within the whole address space. Concentration is similar to spatial locality.

Virtual memory systems successfully exploit locality concepts. The main idea behind virtual memory is that the combined size of the program is allowed to exceed the total amount of physical memory available in the system. The operating system keeps the most actively used pages of the program in main memory, and the rest of the program on secondary storage devices. Most programs at any given time reference only a limited number of their pages; in other words, they exhibit locality of reference, thus making virtual memory system an efficient solution for improving the performance of computer systems.

Cache memory systems in processors exploit locality of reference by storing recently accessed data or instructions. Information can be fetched or pre-fetched into a processor cache depending on the locality characteristics of the reference pattern. By storing frequently accessed information in a processor cache the average time needed to access data and instructions is greatly reduced, since the time required to access the cache is significantly lower than the time needed to access main memory.

Locality of reference is also used in file systems to increase the efficiency of file access. This is similar to processor cache memory systems, only this time blocks of a file that resides on disk are cached in memory, thus reducing the time required for file access.

A number of methods of exploiting locality concepts for improving the performance of network interconnection devices have been proposed. These methods rely on caching references to network addresses, because a sequence of packets on the network can be viewed as a sequence of references. The research presented in this report is concerned with efficient caching techniques that exploit the locality of destination address references in network traffic.

4 Locality in IP Traffic

Several previous studies have identified the presence of locality in the network traffic references. However, few studies have been concerned with locality in large internetworks whose main addressing scheme is regulated by IP (Internet Protocol). This section presents findings about temporal and spatial locality of reference in large IP-based internetworks.

First, locality concepts are adapted to locality in IP traffic. Then experimental methods used to identify locality are presented along with the experimental data. The analysis of temporal and spatial locality is given and some conclusions are presented.

As stated previously, network traffic can be viewed as a sequence of packets where each packet has its own destination address. This sequence of addresses posseses locality of reference. As a result, terms such as temporal and spatial locality can be adapted to the discussion of locality in IP traffic.

Two different types of locality - temporal and spatial - can easily be identified. Temporal locality means that there is a high probability of referencing the same address within a short period of time. That is, an address that has been referenced recently is more likely to be referenced again than one that has not been seen for awhile. This due, in part, to traffic passing across the network in "trains" of packets [Jain 90]. Spatial locality means that there is a high probability of referencing addresses in the same numerical range, or network (*neighborhood*). The region or *neighborhood* may be a group of addresses, a subnet, or a group of subnets.

Arrival Time:	(a)	(b)	(c)	(d)
0.01	192.168.205.76	192.168.205.76	192.168.45.125	192.168.45.125
0.02	192.168.205.76	192.168.205.98	192.168.45.129	192.168.45.127
0.03	192.168.205.76	192.168.205.98	192.168.205.98	192.168.45.126
0.04	192.168.205.76	192.168.205.76	192.168.45.130	192.168.205.76
0.05	192.168.205.76	192.168.205.98	192.168.201.19	192.168.205.98
0.06	192.168.205.76	192.168.205.98	192.658.45.126	192.168.45.125
0.07	192.168.205.76	192.168.205.76	192.168.201.45	192.168.45.127
0.08	192.168.205.76	192.168.205.98	192.168.205.98	192.168.45.126
0.09	192.168.205.76	192.168.205.98	192.168.45.125	192.168.205.76
0.10	192.168.205.76	192.168.205.76	192.168.45.126	192.168.205.98
0.11	192.168.205.76	192.168.205.98	192.168.45.130	192.168.45.125
0.12	192.168.205.76	192.168.205.98	192.168.45.129	192.168.45.127
0.13	192.168.205.76	192.168.205.76	192.168.201.45	192.168.45.126
0.14	192.168.205.76	192.168.205.98	192.168.201.19	192.168.205.76
0.15	192.168.205.76	192.168.205.98	192.168.205.98	192.168.205.98

Figure 1: Example of destination address traces.

4.1 Temporal Locality

Internetwork traffic often consists of sequences of packets that share a destination address; in IPv4, each packet has a 32-bit destination address [Maufer 99]. For example, the sequences of packets depicted in Figure 1 (a) and (b) have a high degree of temporal locality because they both have a high probability of referencing the same IP address in a short period of time. The sequence in Figure 1 (a) consists of 15 references to only one IP address `192.168.205.76` with an interarrival time of 0.01 seconds. This means that 100% of all packets in the trace have interarrival time of 0.01 seconds. However, for the trace in Figure 1 (b) the number of packets with interarrival time of 0.01 seconds is smaller because this sequence of packets consists of two different flows of packets destined to IP addresses `192.168.205.76` and `192.168.205.98` respectively and packets with the same IP addresses are non-uniformly distributed in the trace. 38% of all packets have interarrival time of 0.01 seconds, 31% have interarrival time of 0.02 seconds, and the remaining 31% have interarrival time of 0.03 seconds. This means that the trace in Figure 1 (a) has a higher degree of temporal locality than the trace in Figure 1 (b).

In comparison to the traces in Figures 1 (a) and (b), the sequences in Figures 1 (c) and (d) have relatively poor temporal locality. Trace in Figure 1 (c) consists of 8 flows with each flow destined to a different IP address. The distribution of packets according to their interarrival times for this trace is as follows: 20% have interarrival time of 0.04, another 20% have interarrival time of 0.05, 10% have interarrival time 0.06, 20% have interarrival time of 0.07, 10% have interarrival time of 0.08, 10% have interarrival time of 0.09, 10% have interarrival time of 0.10. Trace depicted in Figure 1 (d) has 5 flows where 50% have interarrival times of 0.05 and another 50% have interarrival time of 0.06.

Based on the above discussion, the definition for temporal locality in IP traffic is stated as follows: a trace of references to IP addresses has high temporal locality when a large portion of the repeated references have a short interarrival time, i.e. a trace has a large probability of re-referencing the same IP address in a short period of time.

4.2 Spatial Locality

To define spatial locality for IP traffic, first the term *neighborhood* should be explained in relation to IP addressing. A neighborhood is the collection of addresses which are close physically or numerically to some given address. For example, the neighborhood for the address `7` is a collection of addresses `{5, 6, _ , 8, 9}`. Neighborhood addresses can easily be observed in memory or on disk. For network addresses a neighborhood is a set of hosts physically close to each other. Then there is no neighborhood relation implied by

the numerical proximity of network addresses, because hosts can be in the same physical location, but on different local area networks that use different addressing schemes. In the case of IP addressing, addresses can be considered in the same group (neighborhood or cluster) if they have a common prefix of some predefined length. Generally, this prefix will be an address of a physical or virtual network on which the hosts are located.

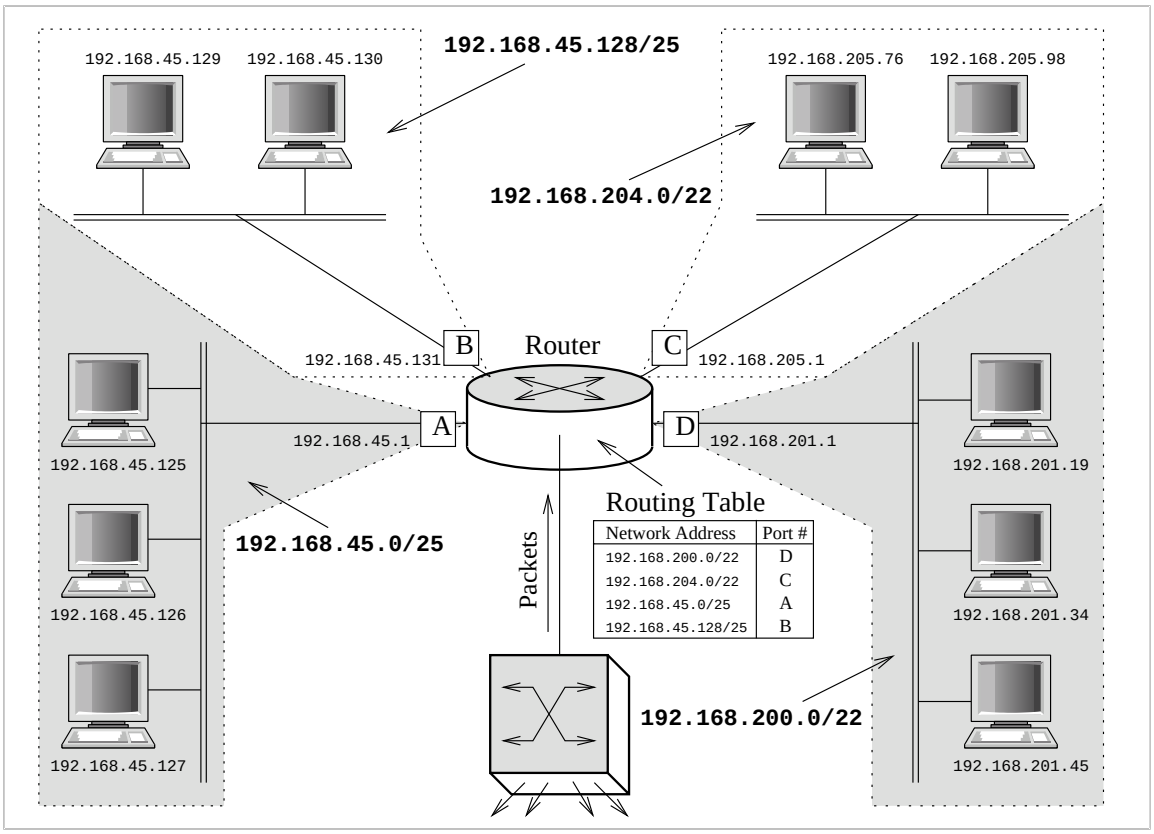

Figure 2: Sample network.

Consider sample network presented in Figure 2. There are 4 subnets which were created using IP addressing [Fuller et al 93]. There is a router in the middle of the network which connects these four networks. The routing table of the router is also shown. Each of these subnetworks can be identified as a neighborhood or a cluster of IP addresses. Now consider the traces presented in Figure 1. Traces (a) and (b) have high spatial locality, because they contain references to only one neighbourhood that is aggregated under the network address of `192.168.204.0/22` . All packets in these traces will be routed, according to the routing table shown, to the subnet connected to the router via port `#C`.

In contrast, the trace presented in Figure 1 (c) has very low spatial locality, because it contains references to all 4 subnetworks of the sample network given in Figure 2: `192.168.200.0/22`, `192.168.204.0/22`, `192.168.45.0/25`, `192.168.45.128/25`. As a result, the packets will

be routed to all 4 subnets that are connected to the router. It contains references to a large number of distinct subnets or clusters in comparison to the total number of subnets available (in this case only 4 subnets are available to be referenced). It should be noted that this particular trace also has very low temporal locality. Thus, it is possible for the sequence of references to have low locality in both dimensions – temporal and spatial.

However, even though the trace presented in Figure 1 (d) has low temporal locality it has higher spatial locality than the trace depicted in Figure 1 (c). This trace has references to only two subnets out of the four possible: `192.168.204.0/22`, `192.168.45.0/25`. According to the routing table some packets in this trace will be routed to subnet connected via port `#C` and other packets to subnet connected via port `#A`. This shows that there is a possibility for traces to have low temporal and high spatial locality.

To summarize, spatial locality for IP traffic is defined as follows: a trace of references to IP addresses has high spatial locality if it references a limited number of the available subnets, i.e. a trace has a high probability of referencing addresses from the same subnet.

4.3 Summary

There is the possibility of traffic locality of both types – temporal and spatial – in internetwork traffic, due to the nature of IP routing, addressing and aggregation schemes. However, one must carefully investigate actual IP traffic in order to implement efficient locality-based caching and routing techniques. Since currently available models do not replicate real IP traffic very well, trace driven simulation is used in this work.

5 Methodology

To analyze different types of locality in IP traffic, several methods were developed based on well-known sampling techniques described in the literature [Claffy et al 93], [Rueda et al 96]. This section describes the methods used to gather information on temporal and spatial locality in internetwork traffic along with the characteristics of the experimental data.

5.1 Experimental Model for Temporal Locality

To assess the degree of temporal locality in a stream of internetwork addresses the distribution of interarrival times of packets was analyzed. The method used to collect the necessary data was based on computing the number of IP destination addresses repeated within a

given time interval. In other words, when characterizing temporal locality the intervals between repeated references to the same address were taken into account rather than just the interarrival time between two successive packets on a link. The count of the number of repeated references was calculated for a range of values of the interval. This was done by scanning the packet trace for the destination address and arrival time of each packet. The interarrival time was computed for each repeated address discovered in the trace, and the number of arrivals in the time bin for that interarrival time was incremented. After the whole trace had been scanned, each time bin contained a count of repeated references that had the corresponding interarrival time.

5.2 Experimental Model for Spatial Locality

For analysis of spatial locality a clustering algorithm was developed. It organizes the address space that was accessed by the whole trace into clusters in order to identify spatial behaviour of the reference stream. The clusters produced by this algorithm are groups of IP addresses with a common prefix. Thus, each cluster can be thought of as representing a subnet with an address mask equal in length to the number of bits in the common prefix of the addresses in this cluster. For example, IP addresses `192.168.45.67`, `192.168.45.101` and `192.168.45.7` can be considered as lying in the same cluster with prefix `192.168.45`, or equivalently in the subnet `192.168.45.0/24`. The input data for the clustering algorithm is the set of unique IP addresses found in a trace.

The algorithm works as follows. First, the starting mask length is selected; the default is a starting mask of zero. Then the set of unique IP addresses is scanned and the mask is applied to every one of them. For example, the address `192.168.45.67` under a mask of length 24 becomes `192.168.45.0` because each of the four decimal numbers in the address represents an 8-bit quantity. The number of IP addresses with the same prefix after applying the mask is counted, and these addresses are recognized as forming a cluster. If there are more than a predefined number (this predefined number is called *cluster size*; for this research cluster size was set at 32) of addresses in a cluster it is split further in the next iteration. If the number of addresses placed in a cluster is less than the cluster size, this cluster is considered complete and all addresses belonging to it are marked as *clustered*. The mask length is then incremented and the procedure is repeated, skipping all IP addresses marked as clustered. The process of declaring clusters when they are less than the limiting value in size is appropriate because it starts with a mask length of zero which potentially covers all addresses in the trace, and iterate towards longer, more specific masks. The number of clusters discovered can be interpreted as the number of subnets present in the original trace.

5.3 Experimental Data

The experimental data used in this research consisted of traces – sequences of IP packet headers – collected at real gateway routers at various sites [WAND], [NLANR], [UA]. The characteristics of the traces discussed in this research are presented in Table 1.

The two traces, *A-1* and *A-2*, are collections of long GPS-synchronized IP headers and corresponding timestamps that were captured with a DAG2 system at the University of Auckland Internet uplink by the WAND research group in November 1999 [WAND]. The tap was installed on an OC3 link (155.52 Mbps) carrying a number of Classical-IP-over-ATM, LANE and POTS services. The trace contains all Classical-IP headers of a single VPI/VCI pair, which connects the university to the local service provider. A maximum 2 Mbps peak packet rate was set in each direction. These traces (*A-1* and *A-2*) have been sanitized by mapping the addresses onto `10.X.X.X` network to preserve privacy. This makes the analysis of spatial locality in these traces impossible. However, temporal locality should be unaffected by the sanitation process. Only ICMP, TCP and UDP packets appear in the trace. For UDP packets and IP fragments all user payload is zeroed. There was no information provided on the type or characteristics of the router that handles the described connection.

Trace	A-1	A-2	SDSC-1	SDSC-2	UofA
Number of IP packets	3626938	30270778	3619341	31518464	999990
Number of unique IP addresses	1622	68167	28474	130163	5797
Run length (estimated)	03:11:30 or 11489 sec	38:29:11 or 138551 sec	00:19:58 or 1198 sec	02:43:38 or 9818 sec	00:01:11 or 71 sec
Packets/second	316	218	3021	3210	141000

Table 1: Traces of IP addresses collected at different sites.

The next two traces, *SDSC-1* and *SDSC-2*, are collections of selected fields of IP packet headers and the timestamps captured at the San Diego Supercomputer Center commodity connection in 1995. The traces were not sanitized which makes them ideal for investigation of both temporal and spatial locality. As with the Auckland traces the router characteristics were not given.

The last trace, *UofA*, was collected in May 2001 at the University of Alberta, at the major connection between the university and the local ISP, Telus Communications Inc.. The trace is a collection of IP packet headers with timestamps obtained by using `tcpdump`. The router responsible for connection between the university and ISP is a Cisco 7507 with R5000 RISC processor at 200Mhz, 256MB of DRAM, 2MB of SRAM for packet buffers and

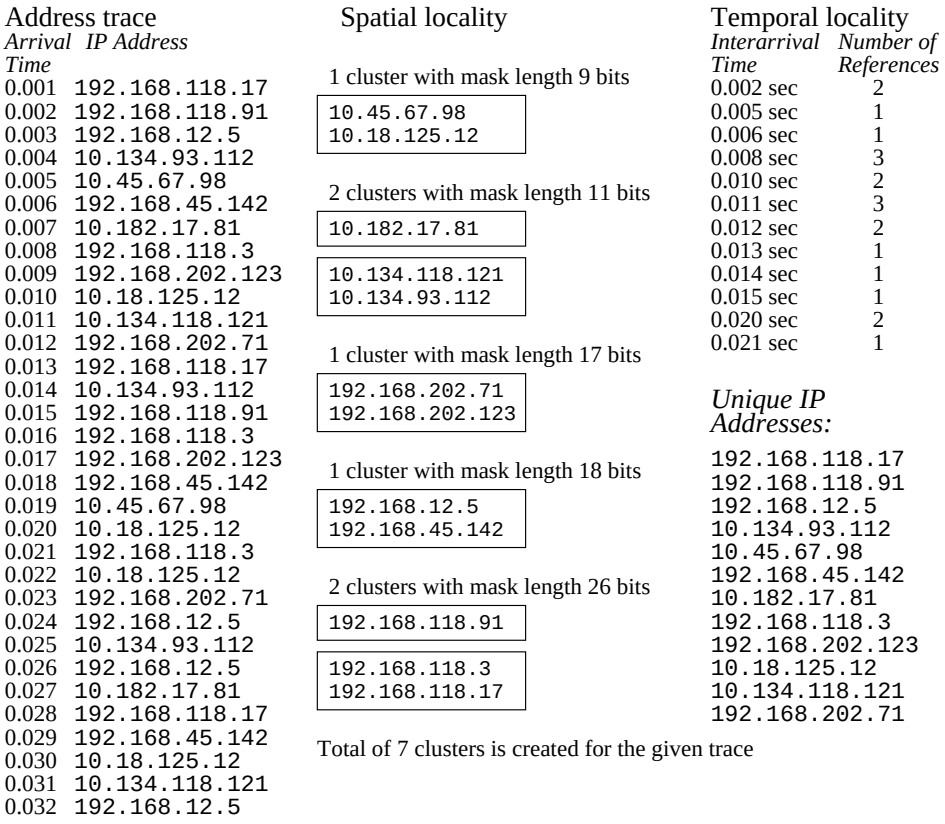

Figure 3: Sample trace for validation of experimental models.

512KB of L2 cache. It has 2Gbps backplane and has seven modular slots. The interface the trace was captured on is 100BaseFx, which is contained in Versatile Interface Processor module which keeps its own copy of forwarding table and does distributed switching. Trace *UofA* was not sanitized to preserve the spatial locality data.

Also, it should be noted that the above traces have different traffic intensity. *UofA* trace has the highest throughput of packets per second in comparison to other traces. Possible reason for this is that the IP traffic intensity on the interface where *SDSC-1* and *SDSC-2* traces were collected is lower than on the interface at the University of Alberta. Detailed comparison could not be done due to the lack of information on the router and interface characteristics at San Diego Supercomputer Center.

5.4 Validation of Experimental Models

The experimental methods described above were implemented using `C` programming language and its standard libraries. To make sure that temporal and spatial locality data collected are valid these models are checked for correctness. Usual approach to this task

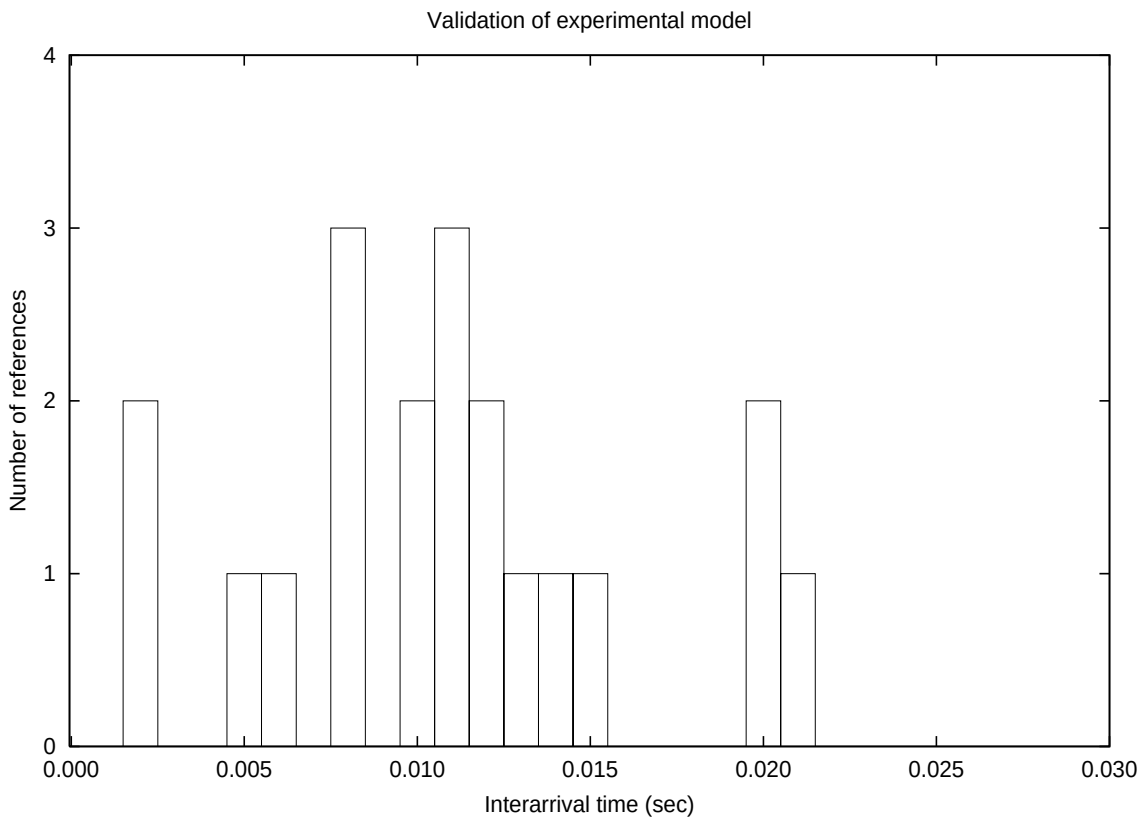

Figure 4: Validation of experimental model for temporal locality.

is to run the experimental models using data sets with known characteristics. Figure 3 presents IP address trace with arrival time of the corresponding packets. Temporal and spatial locality characteristics are shown for the trace.

The experimental models were run with IP address trace presented in Figure 3. Temporal and spatial locality data collected by these experiments match the expected results and are presented in Figures 4 and 5. From the analysis of the data collected follows that experimental models are correct and can be used in analysis of temporal and spatial locality in internetwork traffic.

6 Experiments and Analysis

Measurements of temporal and spatial locality were conducted using the methods and IP address traces described above. For temporal locality experiments the width of each time bin was selected to be 0.001 seconds to ensure more accurate capture of the temporal behaviour of the traces. The range of interarrival times of repeated references to the same IP

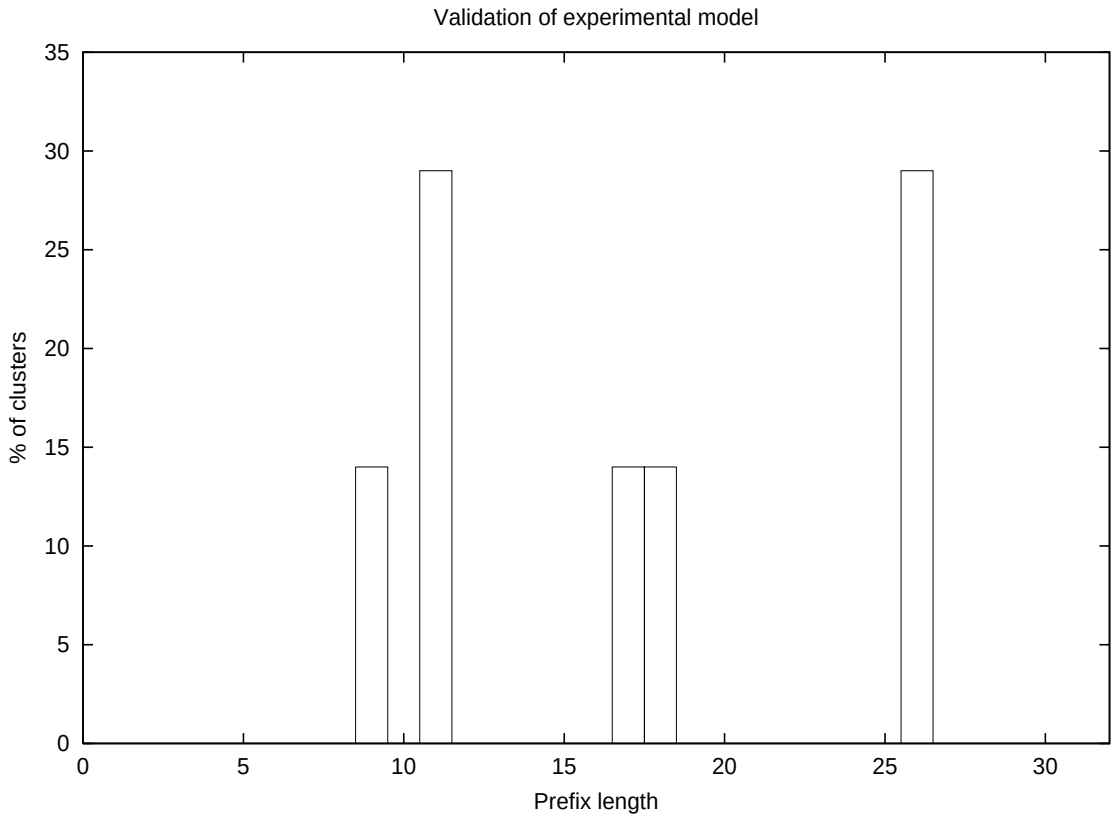

Figure 5: Validation of experimental model for spatial locality.

address was chosen to be between 0 and 5 seconds, which resulted in a total of 5000 time bins. Interarrival times of repeated references to the same address greater than 5 seconds were discarded. These values amounted to 1.7% of the total number of references in the *A-1* trace, 1.4% in trace *A-2*, 2.5% in *SDSC-1*, 3.4% in *SDSC-2*, and 0.0005% in the *UofA* trace. For the spatial locality experiments, the cluster size limit was selected to be 32, i.e. if 32 or less IP addresses form a cluster it is considered to be a complete cluster and it is not split during the later iterations of the clustering algorithm.

Data collected on temporal locality are presented in Figures 6 through 9, and Figures 14 and 15. Spatial locality data are presented in Table 2, Figures 10 through 13 and Figures 16 and 17. Traces *A-1* and *A-2* were not included in the discussion of spatial locality, because their spatial locality was destroyed by sanitation process.

6.1 Temporal Locality

Figures 6 through 8 and Figures 14 and 15 show the distributions of interarrival times of repeated references. Figure 9 presents a comparison of the cumulative distributions for all

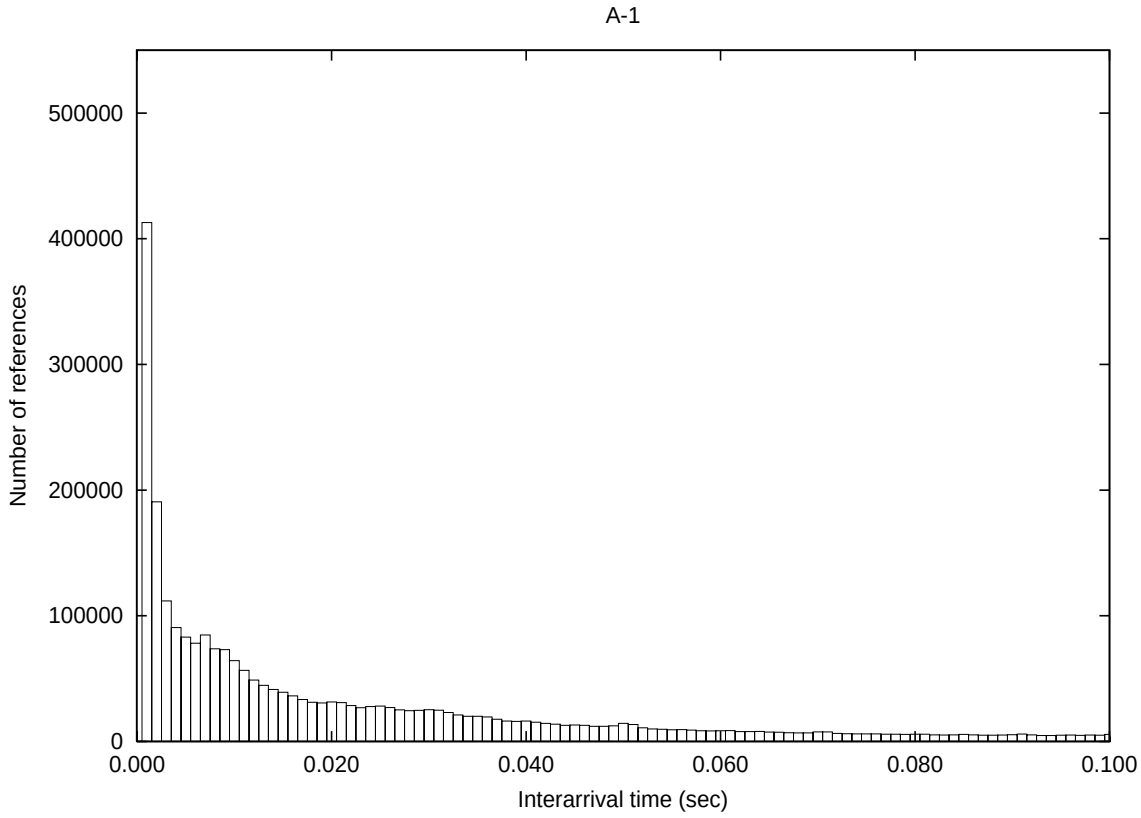

Figure 6: Number of references with interarrival times from 0.001 to 0.100 sec. Trace *A-1*.

traces.

Figures 6 through 9, 14, and 15 show that the distribution of references to the same IP address is non-uniform for all traces, and that a large number of repeated references have very short interarrival times. Approximately 80% of all references have interarrival times less than 0.200 seconds. This percentage is highest for the fastest connection *UofA*. This suggests a high degree of temporal locality, because there is a high probability of referencing the same IP address within a short period of time.

For trace *A-1* 63.9% of all references have interarrival times less than 0.050 seconds and 72.5% occur in less than 0.100 seconds. *SDSC-1* has similar distribution where 53.7% of all references have interarrival times a less than 0.050 seconds and 64.2% are less than 0.100 seconds. The *UofA* trace has even larger proportions of references with short interarrival times: 77.1% are less than 0.050 seconds and 84.7% are less than 0.100 seconds. In addition, the percentage of references with interarrival times of 0.001 seconds or less is overwhelming in all traces: *A-1* - 11.6%, *A-2* - 13.6%, *SDSC-1* - 15.4%, *SDSC-2* - 20.1%, *UofA* - 44.8%. This means that a large number of packets with the same destination IP address follow each

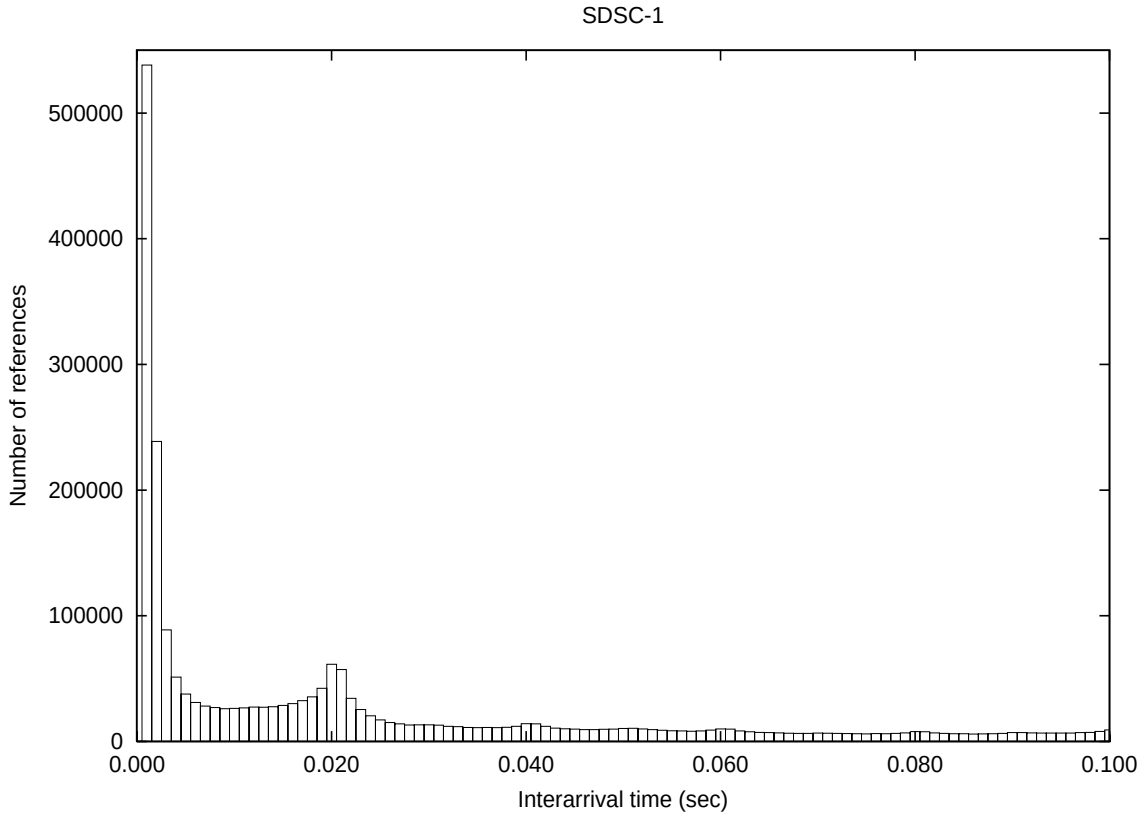

Figure 7: Number of references with interarrival times from 0.001 to 0.100 sec. Trace *SDSC-1*.

other in the shortest time interval recorded. Thus, the probability of referencing the same destination address in a short period of time is very high. This demonstrates that there is a high degree of temporal locality in these traces.

The comparison of cumulative distributions of interarrival times for all traces (Figure 9) shows that large number of all references have very short interarrival times in comparison to the longest recorded. The cumulative distributions in traces *A-1* and *A-2* have interarrival times less than 0.250 seconds for 85.6% and 85.8% of all references respectively. Similarly, traces *SDSC-1* and *SDSC-2* have interarrival times of less than 0.250 seconds for 81.2% and 82.7% of all packets. As for the *UofA* trace, 93.4% of all references have interarrival times less than 0.250 seconds. This provides additional support for the claim that there is a high degree of temporal locality in traces of IP addresses.

The analysis of temporal locality in these traces shows that they have a high degree of temporal locality, thus making it possible to apply the methods and concepts developed for locality of reference in operating systems to the investigation of temporal locality in IP

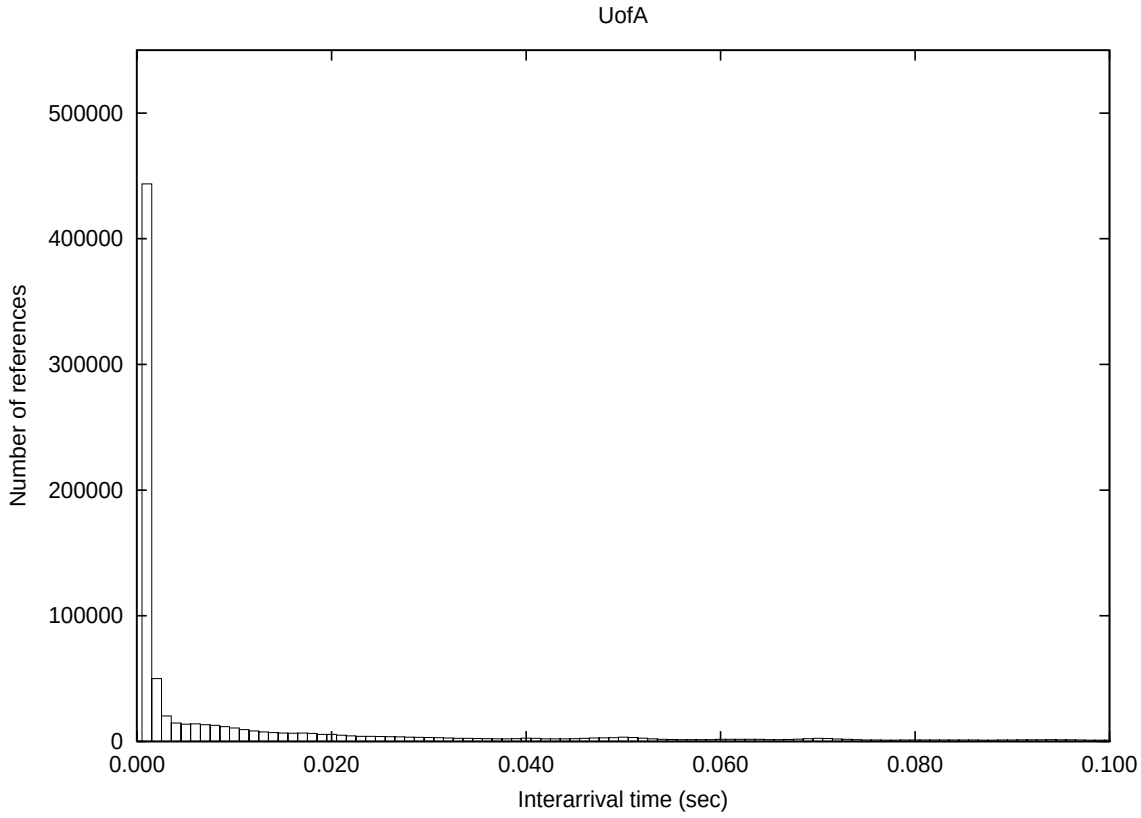

Figure 8: Number of references with interarrival times from 0.001 to 0.100 sec. Trace *UofA*.

traffic.

6.2 Spatial Locality

Table 2 and Figures 10 through 12 present the distributions of clusters in the traces *SDSC-1*, *SDSC-2*, *UofA* according to the prefix length of a cluster. Figure 13 shows the cumulative distribution of clusters for all traces. In this discussion the same definition of cluster is used as in the description of the clustering algorithm.

Figures 16 and 17 present the distributions of clusters in traces *A-1* and *A-2* respectively. When these traces were originally collected, the addresses were mapped onto `10.X.X.X`. This is a common procedure used to preserve privacy when cataloging traces. The clustering results presented in Figures 16 show that addresses in the trace *A-1* are concentrated in clusters with mask lengths between 18 and 26 bits. Figure 17 also shows that in trace *A-2* the clusters are concentrated in one particular region of the address space. The addresses for this trace are concentrated in clusters which have mask lengths between 22 and 27 bits; approximately 92% of all clusters have a mask length of 27 bits. Table 2 presents more de-

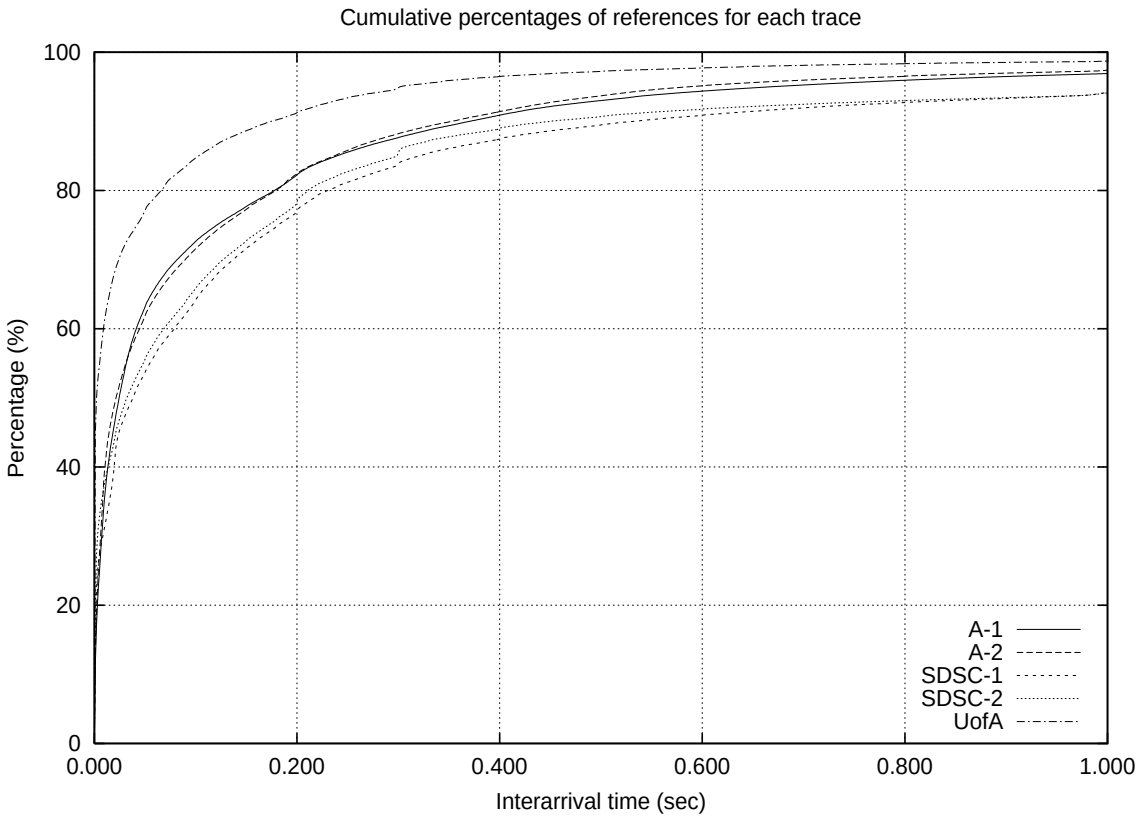

Figure 9: Cumulative percentage of references with interarrival times from 0.001 to 1.000 sec. All traces.

tailed distributions of prefix lengths of clusters. Distributions of clusters in traces *A-1* and *A-2* significantly differ from those of non-sanitized traces. Since mapping of IP addresses in these traces was performed based on location of the IP address in a trace and not on prefix of the address, such mapping assigns IP addresses that have same prefixes different IP addresses in `10.X.X.X` network that do not share same prefix. As a result, packets which are destined to IP addresses on the same subnet are assigned IP addresses that are not related in space and are on different subnets in the address space represented by `10.X.X.X`.

Thus, it is concluded that sanitation process performed on these traces destroyed much of a spatial locality data which makes these two traces (*A-1*, *A-2*) not useful in analysis of spatial locality in IP address stream. Only three traces are considered for further discussion *SDSC-1*, *SDSC-2* and *UofA*, since they were not sanitized or mapped in any way.

The cluster distributions for traces *SDSC-1* and *SDSC-2*, shown in Figures 10 and 11, follow non-uniform bell-shaped patterns. For the trace *SDSC-1*, 79.1% of all clusters have prefix lengths between 12 and 19 bits, and for the trace *SDSC-2* 78.7% of all clusters have

Prefix length	A-1	A-2	SDSC-1	SDSC-2	UofA
0-2	0%	0%	0%	0%	0%
3	0%	0%	0.06%	0%	0.31%
4	0%	0%	0.13%	0%	0.31%
5	0%	0%	0.19%	0.04%	0.62%
6	0%	0%	0.44%	0.14%	2.15%
7	0%	0%	0.38%	0.19%	2.77%
8	0%	0%	0.69%	0.21%	0.92%
9	0%	0%	2.14%	0.19%	4.31%
10	0%	0%	3.97%	0.37%	11.38%
11	0%	0%	3.91%	1.62%	16.31%
12	0%	0%	4.98%	3.06%	16.62%
13	0%	0%	13.93%	3.54%	8.62%
14	0%	0%	20.10%	5.12%	4.62%
15	0%	0%	18.21%	9.58%	2.46%
16	0%	0%	9.39%	16.51%	0.92%
17	0%	0%	4.16%	21.12%	2.15%
18	6.17%	0%	4.03%	14.06%	1.85%
19	16.05%	0%	4.28%	7.27%	2.15%
20	17.28%	0%	2.39%	5.01%	1.85%
21	4.94%	0%	2.02%	3.63%	2.77%
22	14.82%	0.78%	1.32%	1.94%	7.69%
23	14.82%	1.41%	0.63%	1.02%	5.85%
24	14.82%	2.10%	0.95%	1.02%	1.23%
25	8.64%	1.55%	0.69%	0.88%	0.62%
26	2.47%	1.73%	0.38%	0.98%	0.92%
27	0%	92.43%	0.63%	2.50%	0.62%
28-32	0%	0%	0%	0%	0%

Table 2: Distribution of clusters. All traces.

prefix lengths between 14 and 20 bits. There are no clusters with prefixes of 0 to 2 or 27 to 32 bits for trace *SDSC-1*, and for trace *SDSC-2* there are no clusters with prefixes of 0 to 4 bits, nor clusters with prefixes over 28 bits in length.

The distribution of clusters in the *UofA* trace (Figure 12) follows a different distribution. It has two peaks at prefix lengths values of 12 and 22 bits. Overall, 61.9% of all clusters have prefixes of 9 to 14 bits in length and 13.5% have prefix lengths of 22 and 23 bits. There are no clusters whose prefix lengths are 0 to 2 bits, nor clusters with prefixes over 28 bits in length.

The above data suggest that a large number of subnets have prefixes of specific lengths (12 to 19 bits for *SDSC-1*, 14 to 20 bits for *SDSC-2*, and 9 to 14 bits for *UofA*) that are responsible for over 75% of the references in IP traces, which supports the claim that there

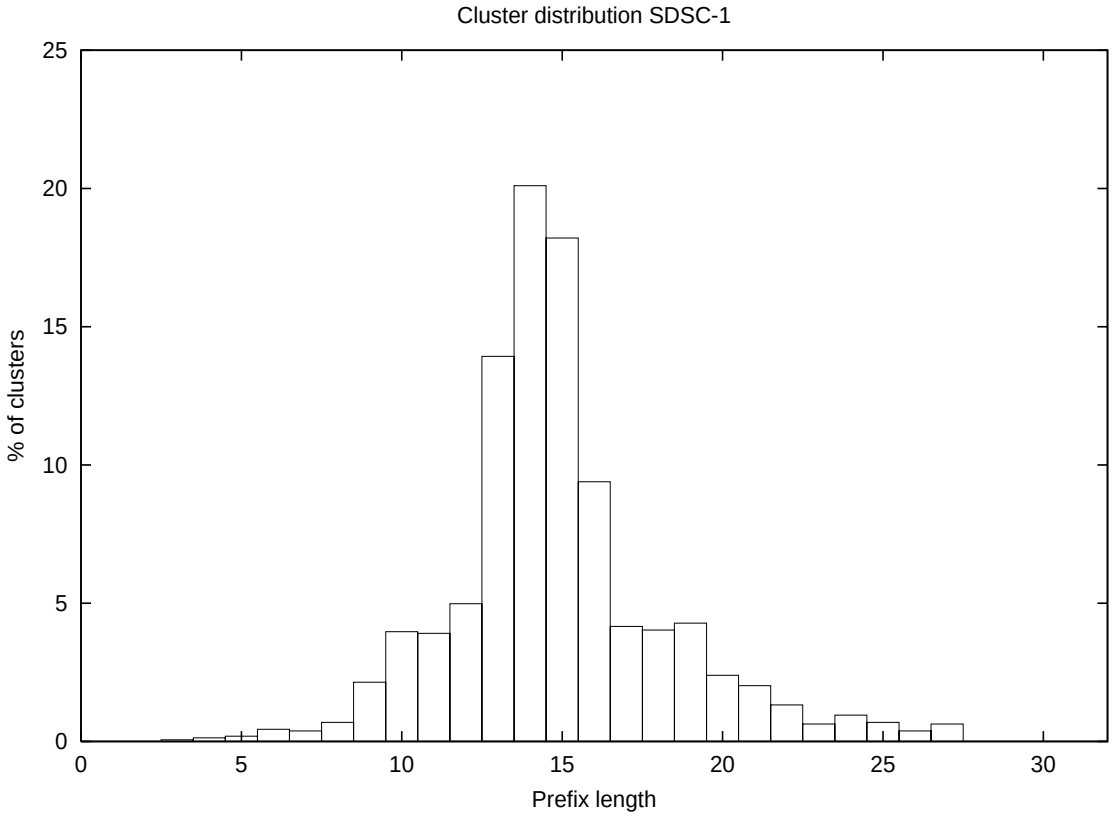

Figure 10: Clustering of *SDSC-1* trace.

is a high degree of spatial locality in IP traffic.

In the cluster distribution for the *SDSC-1* trace (Figure 10), 86.9% of all clusters have prefixes between 10 and 19 bits in length. That is, 86.9% of the clusters are concentrated in 31.3% of the address space. Approximately 95% of all clusters are concentrated in 44% of the whole address space, with prefixes of 9 to 22 bits in length. Similar behaviour is observed for the *SDSC-2* trace (Figure 11). This distribution has 88.9% of all clusters with prefix lengths between 12 to 21 bits, which translates to 88.9% of all referenced network regions being concentrated in 31.3% of the address space. Again, approximately 95% of all referenced clusters are located in 44% of the address space with prefixes between 11 and 24 bits in length. It should be noted that for traces *SDSC-1* and *SDSC-2* the regions with high cluster concentration are contiguous. This supports the conjecture that there is the high degree of spatial locality in IP address traces.

The *UofA* trace has a bimodal cluster distribution with two peaks, where 64.3% of all referenced clusters have prefix lengths between 9 and 15 bits and another 16.3% of all clusters have prefixes between 21 to 23 bits in length. Thus 80.9% of all address regions are

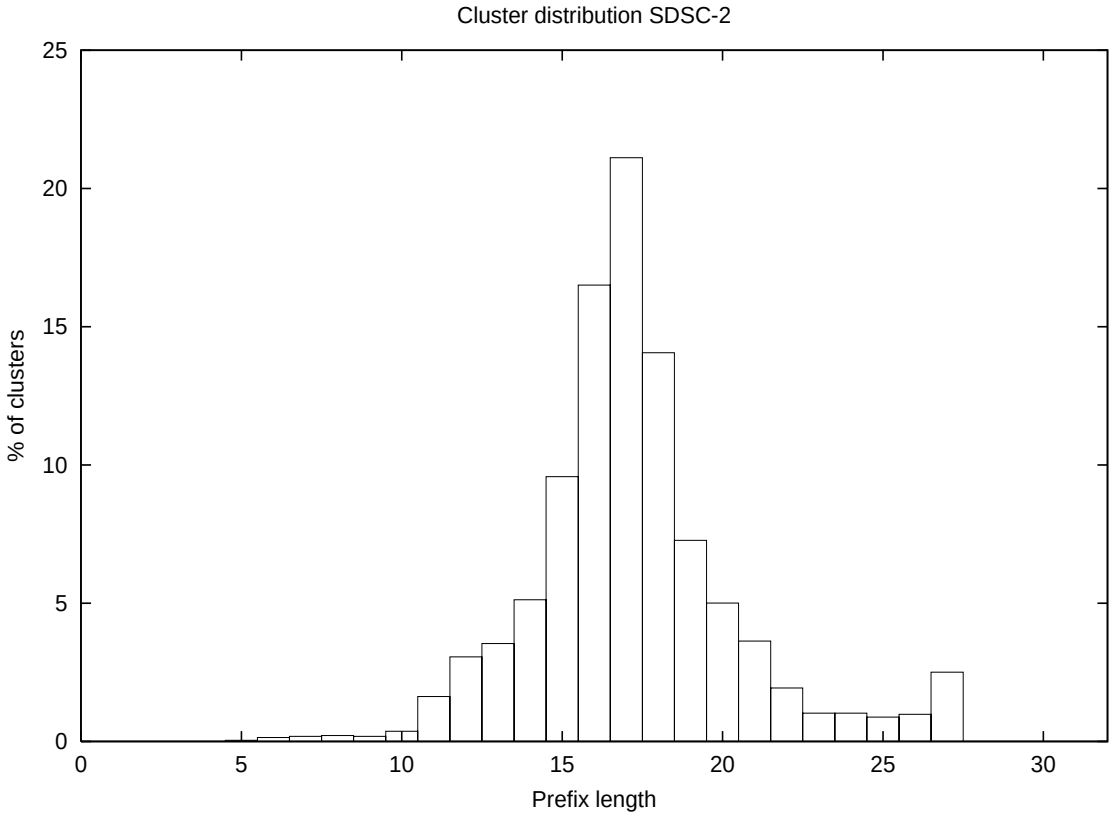

Figure 11: Clustering of *SDSC-2* trace.

concentrated in 31.3% of the address space. As with previous traces, 90.8% of all clusters have prefix lengths of between 9 and 24 bits, so that 90.8% of all referenced clusters are spread over 50% of the address space.

The clusters or subnets that are referenced in IP traces tend to be in a specific region of the address space. Over 90% of the subnets are concentrated in 50% or less of the whole address space available. This provides additional support to the claim that there is a high degree of spatial locality in IP address traces.

Figure 13 depicts cumulative distributions of clusters for all traces. The ranges of prefix lengths that correspond to larger slopes in the graphs are the ranges of high concentration of clusters or subnets. These graphs can be used as a starting point for deriving the parameters that characterize spatial locality in the given traces.

The above data and analysis of spatial locality in IP traffic establish the existence of a high degree of spatial locality in IP traces and provide the basis for exploiting spatial locality in ways similar of those used in operating systems.

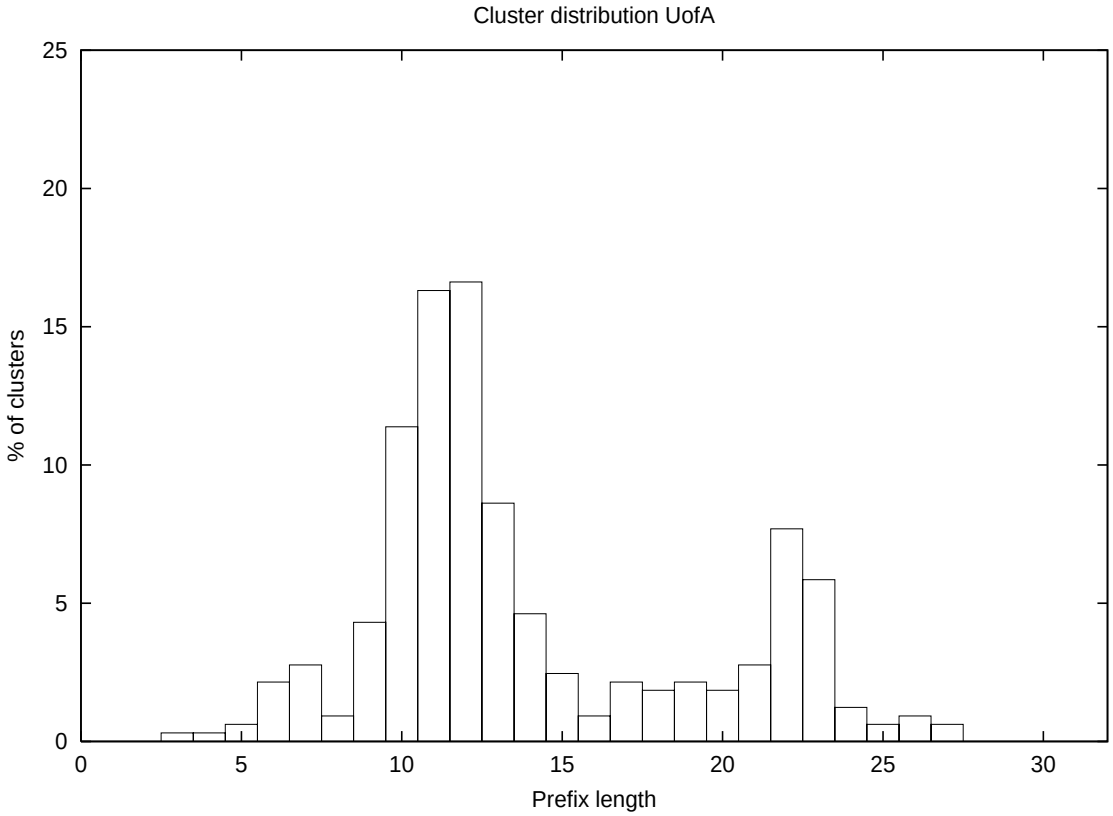

Figure 12: Clustering of *UofA* trace.

6.3 Summary

Several methods were developed in order to collect the required data and analyze the locality behaviour in IP address traces. For the analysis of spatial locality, a new clustering algorithm was developed that captured the distributions of referenced subnets and enabled an analysis of the degree of spatial locality in traces of references to IP addresses. An analysis of temporal locality was based on the interarrival time distribution of references to the same IP address using well-known sampling techniques.

It was found that a large proportion of references has very short interarrival times in comparison to largest recorded – more than 80% of all IP packets have interarrival times less than 0.250 seconds. It was thus concluded that traces of IP addresses have a very high degree of temporal locality.

Data collected on spatial locality in IP traces showed that over 80% of all subnets referenced are concentrated in approximately 30% of the whole address space and that the

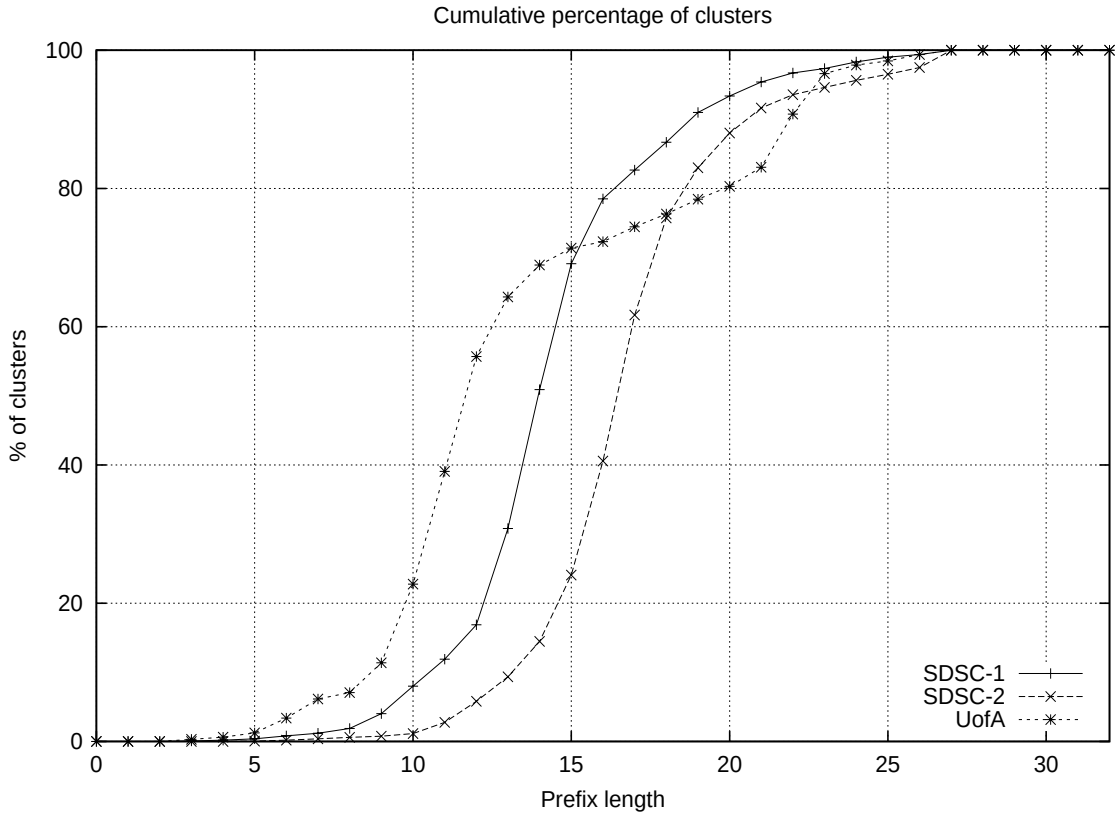

Figure 13: Cumulative percentage of clusters. All traces.

references in IP address traces tend to refer to subnets with specific prefix lengths.

The analysis of traces of IP addresses established the existence of both temporal and spatial locality. One of the suggested applications of locality of reference is in using caches to improve system performance. Locality in IP traffic – temporal as well as spatial – can be exploited in ways similar to virtual memory and file systems to improve the performance of IP address lookup systems.

7 Conclusions

Several methods were developed in order to collect the required data and analyze the locality behaviour in IP address traces. For analysis of spatial locality clustering algorithm was develop that allowed to capture distributions of referenced subnets and to analyze the degree of spatial locality in traces of references to IP addresses. Temporal locality analysis investigated the interarrival time distribution of references to the same IP address using well-known sampling techniques. It was found that large proportion of references has very

short interarrival times in comparison to largest recorded – more than 80% of all IP packets have interarrival times less than 0.250 seconds. It was concluded that traces of IP addresses have a very high degree of temporal locality. Data collected on spatial locality in IP traces showed that over 80% of all subnets referenced are concentrated in approximately 30% of the whole address space and that the references in IP address traces tend to reference to subnets with specific prefix lengths. The analysis of traces of references to IP addresses established the existence of both types of locality behaviour – temporal and spatial. One of the suggested applications of locality of reference is to use caches to improve performance of the systems. The established locality in IP traffic – temporal as well as spatial – can be exploited in a similar way to improve performance of IP address lookup systems, i.e. the introduction of locality-aware IP address cache has potential to significantly improve the performance IP address lookup process.

References

[Belady 66] L.A. Belady, "A Study of Replacement Algorithms for Virtual-Storage Computers", IBM Systems Journal, vol.5, no.2, pp.78–101, 1966

[Bunt et al 84] R.B. Bunt, J.M. Murphy, "The Measurement of Locality and the Behavior of Programs", Computer Journal, vol.27, no.3, pp.238–245, 1984

[Chiueh et al 99 b] T. Chiueh, P. Pradhan. "Cache Memory Design for Network Processors", Submitted for publication to International Symposium on Computer Architecture (ISCA), 1999.

[Claffy et al 93] K.C. Claffy, G.C. Polyzos, H.W. Braun, "Application of Sampling Methodologies to Network Traffic Characterization", Proceedings of ACM SIGCOMM'93, pp. 194–203, 1993

[Claffy 94] K. Claffy, "Internet Traffic Characterization", PhD Dissertation, Department of Computer Science and Engineering, University of California, San Diego, 1994.

[Degermark et al 97] M. Degermark, A. Brodnik, S. Carlson, S. Pink, "Small Forwarding Tables for fast Routing Lookups", Proceedings of ACM SIGCOMM 97, vol.27, no.4, pp.3-14, 1997.

[Denning 70] P.J. Denning, "Virtual Memories", Computing Surveys, vol.2, no.3, pp.153–189, September 1970

[Fuller et al 93] V. Fuller, T. Li, J. Yu, K. Varadhan, "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy", RFC1519, IETF, September 1993

[Hennessy et al 90] J.L. Hennessy, D.A. Patterson, "Computer Architecture: A Quantitative Approach", Morgan Kaufmann Publishers, 1990

[Jain 90] R. Jain, "Characteristics of Destination Address Locality in Computer Networks: A Comparison of Caching Schemes", Computer Networks and ISDN Systems, vol.18, pp.243–254, May 1990.

[Maufer 99] T.A. Maufer, "IP Fundamentals", Prentice Hall PTR, 1999

[NLANR] NLANR Project, San Diego Supercomputer Center, University of California, San Diego; National Science Foundation Cooperative Agreement No. ANI-9807479; and the National Laboratory for Applied Network Research, URL: moat.nlanr.net

[Rueda et al 96] A. Rueda, W. Kinsner, "A Survey of Traffic Characterization Techniques in Telecommunication Networks", Proceedings of IEEE Canadian Conference on Electrical and Computer Engineering 96, vol. II, pp. 830–833, 1996

[UA] Communication Networks Research Group, Department of Computing Science and Computing and Network Services, University of Alberta, URL: www.cs.ualberta.ca/~networks and URL: www.ualberta.ca/CNS

[WAND] WAND Research Group, (Waikato Applied Network Dynamics Research Group), Computer Science Department, University of Waikato, URL: wand.cs.waikato.ac.nz

Data Collected on Traces

This appendix contains data that are discussed, but not presented in the main text of this report.

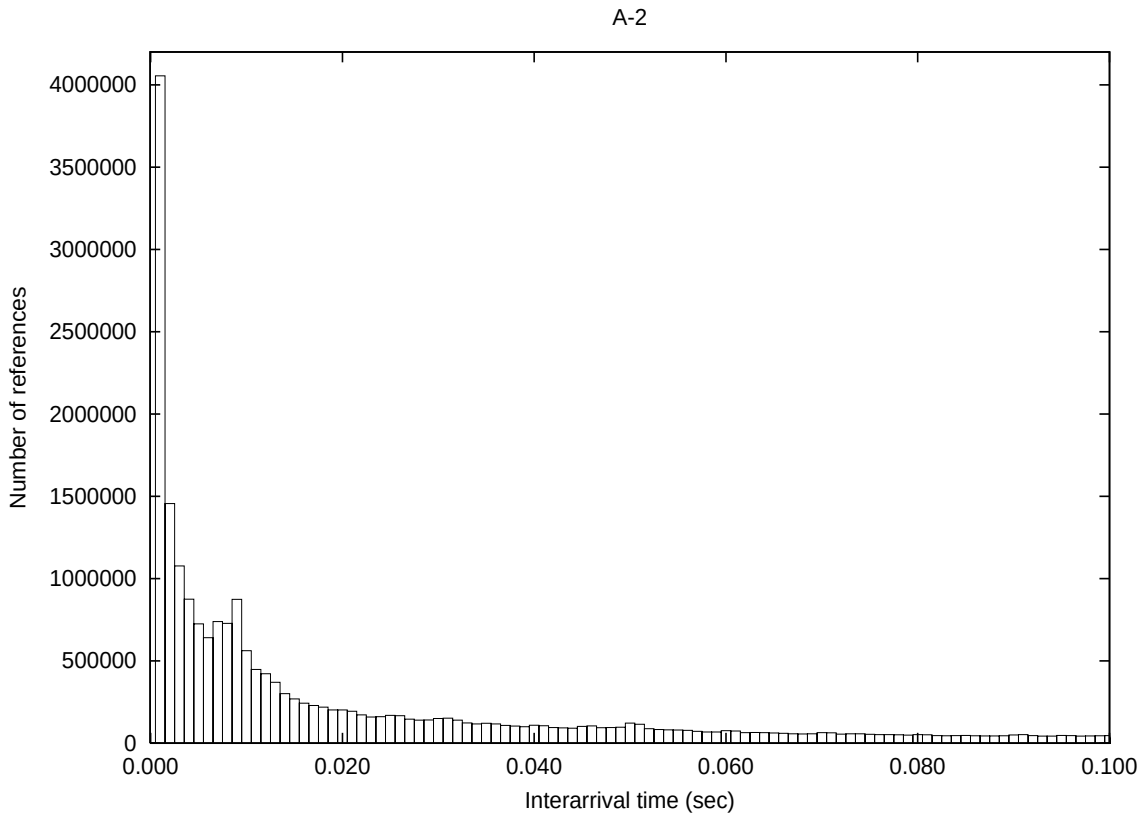

Figure 14: Number of references with interarrival times from 0.001 to 0.100 sec. Trace *A-2*

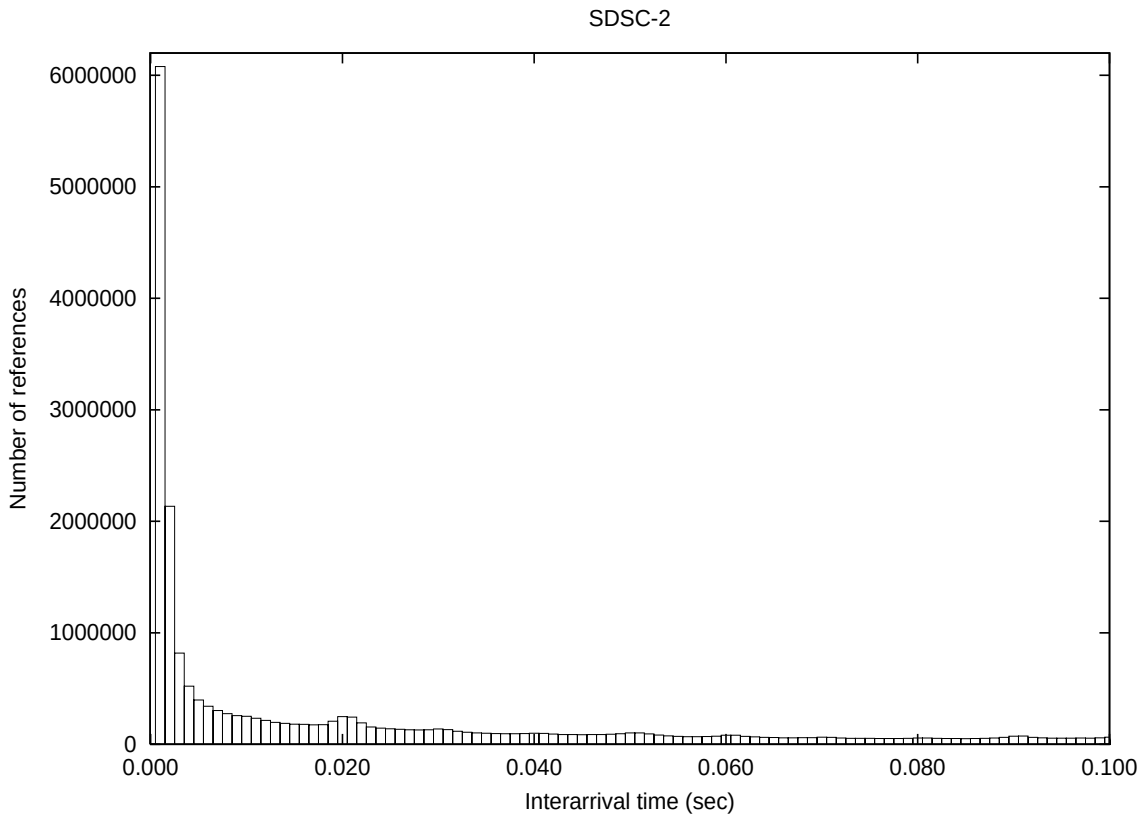

Figure 15: Number of references with interarrival times from 0.001 to 0.100 sec. Trace *SDSC-2*

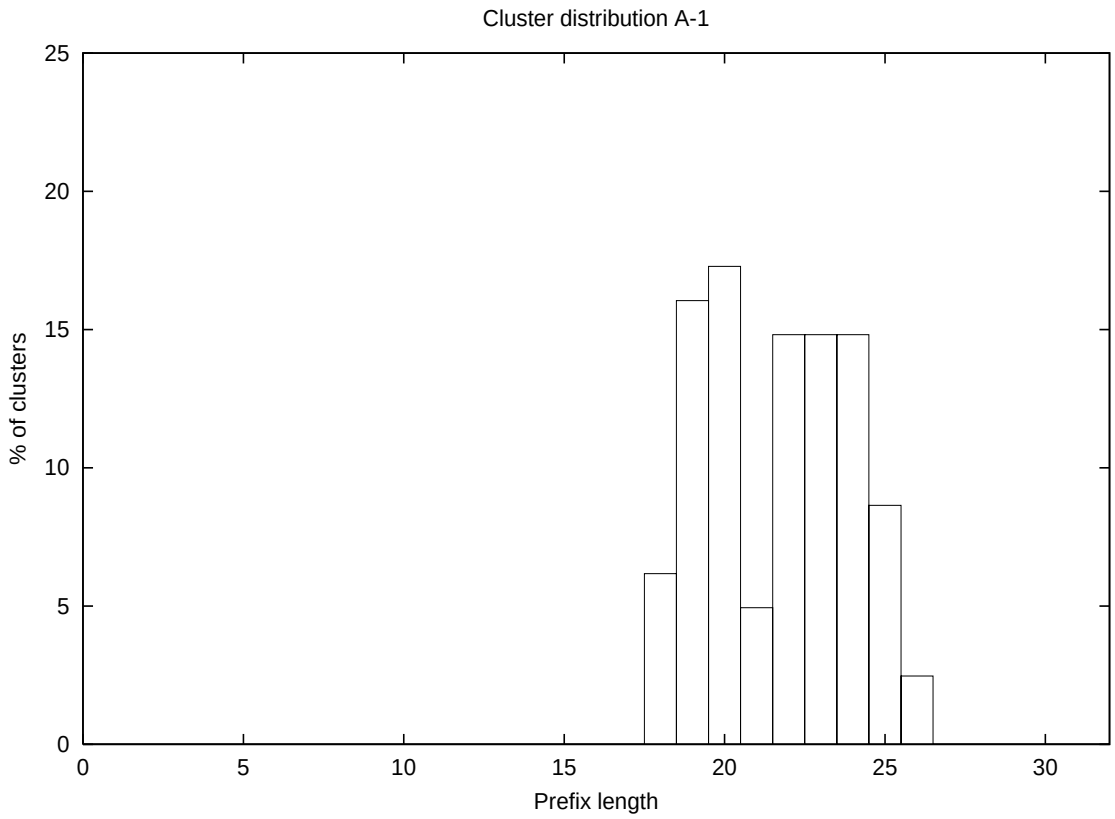

Figure 16: Clustering of *A-1* trace.

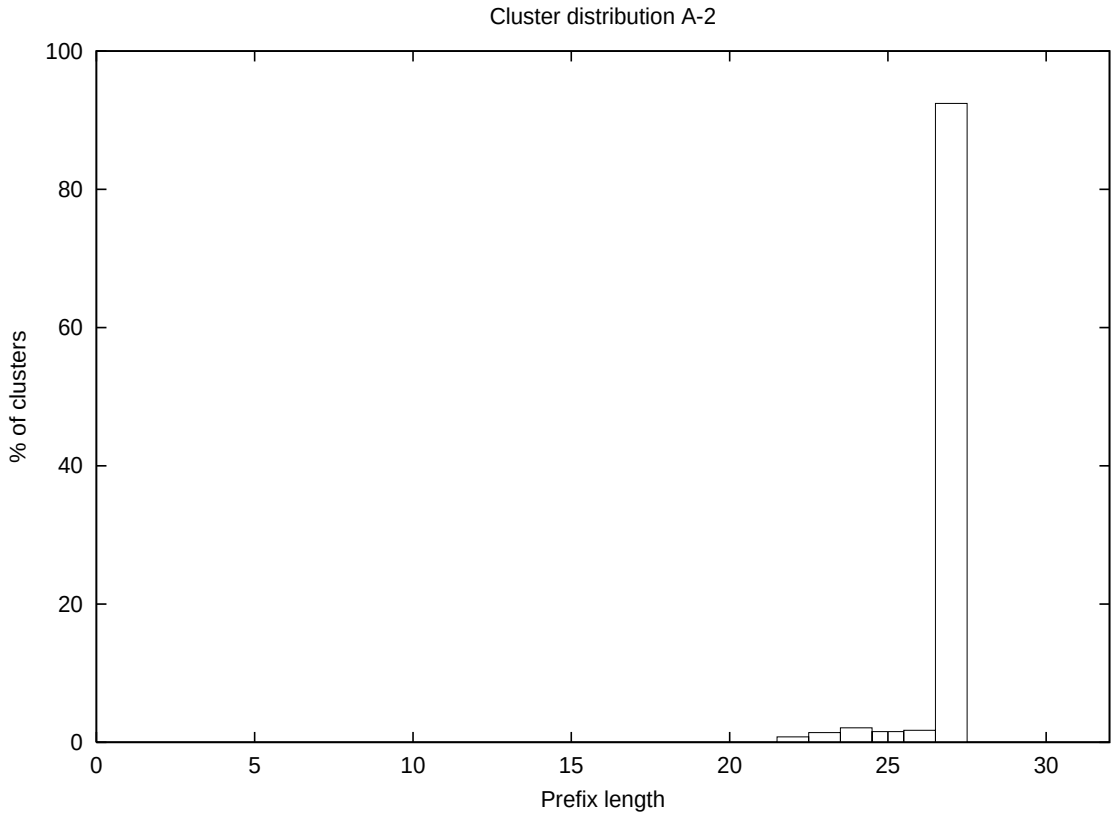

Figure 17: Clustering of *A-2* trace.