An Aggregation Technique for Traffic Monitoring

Kenjiro Cho Ryo Kaizaki Akira Kato
WIDE Project
{kjc,kaizaki,kato}@wide.ad.jp

Abstract

This paper presents an aggregation technique targeted for near real-time, long-term, and wide-area traffic monitoring. Our technique, called aguri, adapts itself to spatial traffic distribution by aggregating small volume flows into aggregates, and achieves temporal aggregation by creating a summary of summaries applying the same algorithm to its outputs. A set of scripts are used for archiving and visualizing summaries in different time scales.

For near real-time monitoring, our prototype implementation employs a Patricia tree and a variant of the LRU replacement policy to limit memory use and search time with variable length keys. The algorithm is fairy insensitive to parameter settings and network conditions.

Aguri does not need a predefined rule set and is capable of detecting an unexpected increase of unknown protocols or DoS attacks, which considerably simplifies the task of network monitoring. We have been monitoring the WIDE backbone network using aguri, and found it useful for network operation.

1 Introduction

Flow-based traffic profiling in which packets are categorized into traffic types and usage information is recorded for each type is commonly used for traffic monitoring [2, 7]. Flow-based traffic monitoring, combined with visualization techniques, provides a powerful tool to understand network conditions [1, 11, 14, 15].

However, a weakness common to the existing flow-based monitoring tools is that, to identify traffic types, predefined filter rules are needed. Filter rules are used to classify packets by examining fields in the packet header. Thus, without *a priori* definitions of traffic types, packets cannot be identified. Flow-based monitoring is facing a difficulty identifying new protocols and dynamically assigned ports. Even for known traffic types, it is not practical to list all possible combinations in the rule set so that minor traffic types are often left undefined and remain unidentified.

On the other hand, the current Internet is exposed to the menace of Denial of Service (DoS) attacks, and DoS attack detection is the highest priority for network operation. The rule-based approach lacks an ability to detect DoS attacks since forged packets can have arbitrary traffic types.

We have been monitoring the WIDE research backbone for years [6], and badly in need of an adaptive monitoring tool for trouble detection, usage reporting and long-term trend analysis. Our focus is traffic measurement to aid network operation, and thus, concise and timely summary reports are more important than precise and detailed reports.

To this end, we have developed a software package called aguri [5]. Aguri uses a traffic profiling technique in which records are maintained in a prefix-based tree and a compact summary is produced by aggregating entries.

Powerful is the feature to produce a summary of summaries applying the same algorithm to its own outputs. Thus, derivative summaries can be produced in different time scales desirable for a specific monitoring purpose. A set of scripts have been developed to visualize summaries. It is also possible to extend the profiler as a protective measure against DoS attacks by using its outputs for traffic control.

Aguri is targeted for near real-time, long-term, and wide-area traffic monitoring. Because automatic aggregation is used for profiling, our approach provides rough usage reports which may not be precise so that it is complementary to the existing tools.

2 Overview

The core idea of an aggregation-based profiling is to aggregate flow entities for profiling. Small volume flows are aggregated until the volume of the aggregate becomes large enough to be identified. A summary output reports the profile of aggregates. An entry in an address profile can be a single host if it consumes a certain portion of the total traffic, or an aggregate if each host entry is small but the aggregate becomes non-negligible. Thus, a limited number of entries are produced, yet it never fails to report high volume entries.

Figure 1 illustrates the concept. A tree before aggrega-

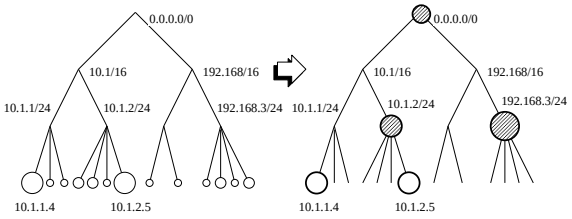

Figure 1. aggregation profiler concept: small entries are aggregated into aggregates

tion is on the left and the corresponding tree after aggregation is on the right. Each node in the tree shows the address space represented by an address prefix and its prefix length. A leaf node corresponds to a single address. The size of a node shows the traffic volume of the node. The usage information recorded at leaf nodes can be aggregated to the shaded internal nodes in the right tree, and a summary reports only the remaining nodes in the right tree.

Summary Profile

It is important to produce concise summary profiles. When a traffic profile is too detailed, important symptoms are buried in excessive data, and often left unnoticed. Each summary profile produced by aguri is compact since small entries are aggregated in a profile.

Aguri produces four separate profiles for source addresses, destination addresses, source protocols and destination protocols. IP addresses are designed to be hierarchical and aggregatable so that it is natural to apply aggregation. Both IPv4 and IPv6 are supported in address profiles. Although protocol numbers are not hierarchical, the same technique can be used to identify port ranges. We concatenate the IP version, the protocol number and the TCP/UDP port number to create a 32-bit key for a protocol profile. A summary reports the total byte count used by each aggregate.

The four separated profiles are effective to capture hostile activities. A victim of a distributed DoS attack will be easily identified in the destination address profile. An originator of port scanning will be identified in the source address profile. A random attack will be identified as a range of addresses as long as some locality exists for the targets. If the locality is unusually low, it is another symptom of a random attack.

Spatial Aggregation

The basic algorithm of the spatial aggregation is quite simple. If there is no resource constraints such as memory consumption or execution time, we could profile every address and protocol occurrence in every packet and, at the end, aggregate entries whose counter value is less than an aggregation threshold. This approach would be acceptable for post-analysis of a saved packet trace. For near real-time monitoring, however, we approximate the above algorithm in exchange for the precision, by managing a fixed number of nodes in the tree using a variant of the Least-Recently-Used (LRU) replacement policy.

When a leaf node is reclaimed, the counter value of the node is aggregated to its parent node. The advantage of this approach is that counter values are never lost even though the resolution is reduced.

To produce a summary output, aguri walks through the tree in the post-order and aggregates nodes if the counter value of a node is less than the aggregation threshold, or outputs the node information if the counter value is above the threshold.

To continue profiling, it is enough to reset the counter of each node; the current tree and the LRU list are kept in tact as a cache, and used for the next profiling period.

Temporal Aggregation

The same algorithm can be used to produce a summary of summaries. Aguri can read its summary outputs, reaggregate them, and produce a new coarse-grained summary. For instance, a 1-hour-long summary can be created out of 60 1-minute-long summaries.

In this paper, an “initial summary” is used to represent a summary directly produced from non-aggregated sources such as captured packets. A “derivative summary” represents a summary produced from summaries.

The method is suitable for archiving profiles since a summary can be created in different time scales from a set of archived summaries. It is also possible to control the resolution by changing the aggregation threshold. The process to generate and archive derivative summaries can be easily automated. Network operators will usually look at only coarse grained summaries but can look into fine grained summaries if necessary.

Archiving and Visualization Utilities

A summary output is in a plain text format so that it is easily processed by various scripts. For archiving, a script is periodically invoked to generate and archive derivative summaries in different time scales such as hourly, daily, monthly, and yearly summaries. The size of a summary is about 5KB so that a small amount of disk space is required for archiving summaries.

Text-based summaries can be converted to a variety of visual images. We have developed a set of scripts for visualization to aid operators to find unusual conditions in summary outputs.

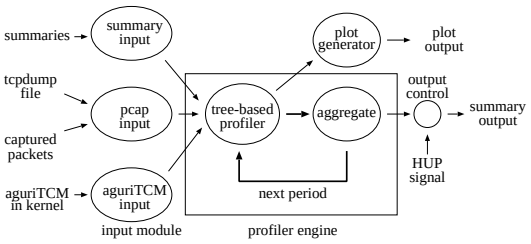

Figure 2. aggregation profiler implementation model

Application for Traffic Control

Once aggregates are identified and profiled, the profile records can be used for traffic control. There are many possible approaches to control aggregates. For example, a rate-limiter can be installed at a firewall to protect the network from a high-bandwidth aggregate [12].

We propose an aguri three color marker (aguriTCM) that combines an aggregation-based profiler with a preferential packet dropping mechanism. The aguriTCM identifies aggregates whose traffic volume is more than the fairshare, and probabilistically raises the drop precedence for those aggregates. The aguriTCM provides rough traffic management based on aggregates in best-effort traffic; the resolution of the control is limited by the resolution of an aggregate in the profile.

Our approach uses Diffserv components as building blocks but the primary target is a stand-alone protection mechanism to minimize the effect of DDoS or flash crowd in best-effort traffic. It also provides rough fairness among aggregates. The details of the traffic control mechanisms are described in [5].

3 Implementation

Aguri, as shown in Figure 2, is implemented as a user program on UNIX. The input modules on the left translate different input formats into a 4-tuple (tree, key, prefix-length, count) and pass them to the profiler engine in the center. Aguri prints summaries to the standard output or a file.

The first input module reads aguri's summary outputs from files or from the standard input to produce a derivative summary. The second input module is an interface to the pcap library [10] that captures packets from a live network or reads a packet trace file saved by tcpdump [9]. The pcap interface allows us to evaluate our prototype using various tcpdump trace files. The third input module reads binary profiles produced by the aguriTCM in the kernel.

```
%!AGURI-1.0
%%StartTime: Sat Jan 06 14:00:00 2001
%%EndTime:   Sat Jan 06 14:00:05 2001
%AvgRate: 17.05Mbps

[dst address] 10658367 (100.00%)
0.0.0.0/0 105652 (0.99%/100.00%)
 0.0.0.0/2 196398 (1.84%/1.84%)
 128.0.0.0/1 141492 (1.33%/97.17%)
        133.28.0.0/16 146217 (1.37%/11.08%)
                133.28.21.21 179320 (1.68%)
         133.28.128.0/17 257220 (2.41%/8.03%)
                133.28.128.14 127541 (1.20%)
                133.28.202.127 470854 (4.42%)
   152.0.0.0/5 157159 (1.47%/25.69%)
        152.10.0.0/16 336636 (3.16%/20.28%)
         152.10.0.0/17 433037 (4.06%/15.16%)
                152.10.1.247 1182481 (11.09%)
                152.10.135.189 208992 (1.96%)
      156.96.0.0/11 253884 (2.38%/3.94%)
        156.114.0.0/16 165979 (1.56%/1.56%)
   168.0.0.0/5 315417 (2.96%/47.96%)
                168.89.12.93 275740 (2.59%)
      173.96.0.0/12 465797 (4.37%/42.42%)
          173.106.176.0/20 248236 (2.33%/38.05%)
                173.106.177.162 440466 (4.13%)
                173.106.177.163 550897 (5.17%)
                173.106.177.172 602230 (5.65%)
                173.106.177.173 1498198 (14.06%)
                173.106.187.134 559784 (5.25%)
                173.106.187.135 155322 (1.46%)
   192.0.0.0/5 111918 (1.05%/8.45%)
    194.0.0.0/7 375630 (3.52%/7.40%)
                194.105.251.45 168327 (1.58%)
                195.130.218.237 244270 (2.29%)
  208.0.0.0/4 283273 (2.66%/2.66%)
%LRU hits: 82.62% (14511/17564)
```

Figure 3. a sample output of a destination address profile

The profiler engine consists of the tree-based profiler and the aggregation module. The tree-based profiler accepts 4-tuples from one of the input modules, and maintains profile records in the trees. At the end of a profiling period, the aggregation module is called to produce a summary. While the aggregation module is walking through the tree in the post-order, each node is either aggregated or reported. To continue profiling, the profiler engine repeats this cycle.

3.1 Summary Output

Figure 3 shows an example of aguri's summary output. A summary starts with a header block, followed by a body block. Lines start with % are comment lines. The body block contains 4 profile types by default but only the destination address profile is shown in the figure. [1]

In the address profile, each row shows an address entry and is indented by the prefix length. The first column shows the address and the prefix length of the entry. When the prefix length is the full length, it is omitted in the output. The second column shows the cumulative byte count. The third

[1] IP addresses appearing in this paper are scrambled for privacy.

```
[ip:proto:srcport] 10570555 (100.00%)
0/0:0:0 4967 (0.05%/100.00%)
4:0/3:0 290382 (2.75%/99.95%)
4:6:0/0 164255 (1.55%/96.15%)
   4:6:0/3 540369 (5.11%/93.38%)
                4:6:20 663178 (6.27%)
                4:6:80 7329218 (69.34%)
         4:6:1024/8 106427 (1.01%/1.01%)
         4:6:1280/8 139741 (1.32%/2.75%)
          4:6:1280/9 150514 (1.42%/1.42%)
        4:6:1536/7 182444 (1.73%/1.73%)
     4:6:2048/5 564594 (5.34%/5.34%)
                4:6:6346 194004 (1.84%)
 4:6:32768/1 128925 (1.22%/1.22%)
                4:17:53 111537 (1.06%)
%LRU hits: 60.80% (10644/17506)

[ip:proto:dstport] 10570555 (100.00%)
0/0:0:0 4967 (0.05%/100.00%)
4:0/3:0 401919 (3.80%/99.95%)
4:6:0/0 579078 (5.48%/96.15%)
          4:6:0/9 327066 (3.09%/4.54%)
                4:6:80 152813 (1.45%)
        4:6:1024/7 419016 (3.96%/17.12%)
          4:6:1024/9 781275 (7.39%/7.39%)
         4:6:1280/8 609679 (5.77%/5.77%)
        4:6:1536/7 597213 (5.65%/12.77%)
         4:6:1536/8 752782 (7.12%/7.12%)
       4:6:2048/6 666539 (6.31%/21.84%)
        4:6:2048/7 155545 (1.47%/15.54%)
          4:6:2176/9 387335 (3.66%/7.96%)
           4:6:2176/10 454168 (4.30%/4.30%)
         4:6:2304/8 645406 (6.11%/6.11%)
       4:6:3072/6 893343 (8.45%/8.45%)
     4:6:4096/4 172569 (1.63%/9.51%)
        4:6:4608/7 688892 (6.52%/6.52%)
                4:6:6346 143558 (1.36%)
  4:6:49152/2 492936 (4.66%/16.44%)
                4:6:49249 1107484 (10.48%)
                4:6:49635 136972 (1.30%)
%LRU hits: 53.96% (9446/17506)
```

Figure 4. a sample output of protocols and ports

column shows the percentages of the entry and its subtree.

The input for this example is a 5-second-long packet trace taken from a trans-pacific link of the WIDE backbone. The parameters of aguri is configured with 256 nodes and 1% aggregation threshold. Among 17,564 observed addresses, only 14 addresses are identified as individual addresses. 38.05% of the traffic belongs to 173.106.176/20; within this address space, 6 distinct addresses are identified. The number of individual addresses found in a typical summary is from 5 to 20. In our trans-pacific profiles, several individual addresses are still identified even in monthly summaries.

A source address profile looks similar. A source address profile tends to identify popular www or ftp servers, whereas a destination address profile tends to identify proxy servers and mirror servers.

Figure 4 shows source and destination protocol profiles. The first column shows a 32-bit key concatenating the IP version number (8bits), the protocol number (8bits), and the TCP/UDP port number (16 bits). For example, "4:6:80" represents IPv4/TCP/HTTP.

In this summary, 96.15% of the total traffic is TCP. Only four individual ports, TCP port 20 (ftp-data), 80 (http), 6346 (gnutella), UDP port 53 (dns), are identified in the source address profile. Note that the use of gnutella is automatically detected without any knowledge about gnutella's use of port 6346.

The destination protocol profile includes a larger number of dynamically assigned ports which are usually aggregated and shown as port ranges. A source protocol profile tends to identify protocols used by servers, and a destination protocol profile tends to identify clients.

3.2 Aggregation Mechanism

The profiler engine implements the prefix-based aggregation algorithm. To produce summaries continuously in near real-time, we need an efficient algorithm in terms of CPU power and memory usage. An approximation limits the number of entries used in a tree, and thus, will make more aggregation than the ideal algorithm. As a result, it introduces two types of errors: (1) part of the counter value could be aggregated to the ancestors, and (2) the entry of a node close to the aggregation threshold could be removed and may not show up in the summary. These errors lower the precision but the impact would be limited. After all, these errors are unavoidable for derivative summaries since aggregation discards details. However, if an entry consumes a non-negligible volume of the total traffic, any approximation will be able to detect it.

To limit memory use and search time with variable length keys, we employ a Patricia tree. Patricia has been employed in the BSD kernel for the internal representation of the routing table [17], and its performance characteristics are well understood. It is suitable to handle 32-bit IPv4 addresses and 128-bit IPv6 addresses.

Patricia is a full binary radix tree. All internal nodes have exactly two children so that when the number of leaf nodes is N, the number of internal nodes is $(N-1)$. Thus, it is suitable for use with a fixed number of nodes, and nodes can be preallocated.

Each node has a prefix as a key associated with its prefix length. The key of an internal node is the common prefix of its two children.

Our use of Patricia is different from the routing table. While the routing table lookup requires best-match, we have only exact-match. In our scheme, a new node is always created when no matching node is found. If there is no available free node, an old node is reclaimed to keep the number of nodes in the tree. Thus, node insertions and deletions occur during a lookup operation.

To update an entry record, the profiler looks up the entry in the tree, and updates the counter value of the entry.

A lookup starts from the root node to a leaf node, checking prefix-matching. If the prefix matches with the internal node, the bit at $(prefixlen + 1)$ of the search key indicates which branch to follow; if the bit value is 0, take the left branch, otherwise, take the right branch. If the matching leaf node is found, the search terminates and the counter of the node is updated.

If the prefix does not match, it indicates no matching node exists in the tree. A new node is created and inserted into the tree. The key is assigned to the new node, and the count is set to the counter. An insertion always creates a leaf and a branch point since single branching is not allowed. The new branch point is inserted as a parent of the unmatching node; the other child of the branch point is the newly created leaf node. The common prefix of the two child nodes is assigned to the branch point. Similarly, deleting a leaf node removes the leaf and its parent. When deleting a node, the counter value is aggregated to its parent.

A fixed number of nodes are preallocated for a tree, and a variant of the LRU replacement policy is used for managing leaf nodes. If the number of nodes is 256, the tree has 128 leaf nodes since $(N - 1)$ internal nodes are needed for N leaf nodes. The LRU is selected because it is simple, cheap and well-understood. The precision could be further improved by using an elaborate algorithm such as the frequency-based replacement [16] but there is a tradeoff between the precision and the efficiency. As already mentioned, the precision is not so important in our scheme.

Since the LRU reclaims a node even when its counter value is very large, a simple heuristic is added not to reclaim a node if the sum of the counter values of the node and its parent is larger than a threshold. The current reclaim exemption threshold is set to 3.123% or 1/32 of the total count.

In the middle of a profiling period, a snapshot of the tree contains nodes with small count values. Nodes whose count value is less than the aggregation threshold are aggregated at the end of the profiling period. The aggregation threshold is set to 1% of the total count by default. The profiler walks through the tree in the post-order so that aggregation and summary output can be done in one pass.

To continue profiling, the profiler just resets the counters and keeps the tree and the LRU list in tact as a cache for the next profiling period. The profiler could reset the counters when aggregating the nodes. However, a two-pass method is used in the current implementation to show the sum of the subtree for readability. The aguriTCM, on the other hand, omits the subtree sum and employs a one-pass method.

IPv4 and IPv6 addresses have different key length. They could be managed in a single tree but separate trees are currently used for ease of debugging. The aggregation threshold is computed from the combined total count so that there is no difference in the summary. On the other hand, the key length is the same for protocol trees so that the profiler uses merged trees.

The profiler uses the same algorithm to produce derivative summaries but there are subtle differences. The size of input sets is much smaller and there are less constrains on execution time or resource usage. Another difference in the Patricia algorithm is that internal nodes are added to insert aggregates, while only leaf nodes are added for initial summaries. A single implementation is currently used for both initial and derivative summaries to reduce the maintenance cost but it could be separately optimized.

3.3 Evaluation Results

We have evaluated the algorithm using packet traces taken from the WIDE backbone. We briefly review the results in this section but the details are described in [5].

Accuracy

In our algorithm, the resolution of aggregation depends on the aggregation threshold. Excessive aggregation can be introduced by the approximation mechanisms so that the number of nodes used in a tree, the replacement policy, the generation of derivative aggregation also affect the accuracy. Although accuracy is not the most important factor to the algorithm, it is better to understand the impact to the results.

We evaluated our LRU-based algorithm with varying number of nodes and varying profiling period length, with or without the heuristic added to the LRU algorithm. The results are compared with the ideal results in which there is no restriction on the number of nodes.

As expected, the simple LRU works well when there are enough nodes but the distortion becomes larger when nodes are insufficient. The tree size of 128 or 256 works well for our backbone packet traces. The aggregation exemption reduces distortion, especially when the profiler runs out of nodes. Therefore, this heuristic works as a safeguard against node shortage.

The effect of the different period length are tested by the traces with different length. Even though the number of the included addresses differs in orders of magnitude, the results look similar. It suggests that there is a locality in address occurrence, and thus, the results are not affected much by the period length.

We also evaluated the impact of summary generations which have different levels of derivative summaries to produce the final results. The results show that the distortion introduced by summary generations is fairly small, which justifies our approach to create derivative summaries for temporal aggregation.

In summary, the algorithm is fairly insensitive to variations in networks, the profiling period length, and summary generations. The packet traces used for the evaluation are backbone data, and as such, the number of included addresses are considerably larger than enterprise networks. The profiler performs much better in enterprise networks.

Performance

For every packet, aguri looks up the matching entry in the 4 trees and manages the LRU lists. When the number of nodes in a tree is N, the lookup operation runs in $O(\lg N)$ time. On the other hand, the cost of managing the LRU list is independent of the number of nodes and it runs in $O(1)$ time. As the number of nodes in a tree increases, the height of the tree becomes longer and the lookup operation becomes more costly.

Our test result shows that the profiler can process about 250Kpps with 256 nodes, and about 200Kpps with 2048 nodes on a PentiumIII 700MHz. The performance is good enough to monitor a 100Mbps link. In the worst case where a 100Mbps link is filled with 64-byte packets, about 190Kpps is required.

3.4 Archiving and Visualization Utilities

Archiving

Aguri prints summaries to the standard output or a file. On receiving a HUP signal, the output file is reopened so that the output file can be redirected to a new file. To archive summaries, a script is periodically invoked by *cron*. The script saves the current output file and sends a HUP signal to the running aguri program to switch the output file.

In our current setting, aguri produces a new summary every 5-seconds. A new summary file containing 24 summaries is created every 2-minutes. The script also generates hourly/daily/monthly/yearly summaries when crossing the time boundaries. It is also possible to customize the script to detect a certain condition and send an alert to the operator.

A summary output size varies depending on the traffic but is usually about 5KB. Uncompressed derivative summaries take about 150KB/hour, 3.5MB/day, 105MB/month and 1.2GB/year. If the initial summaries created every 5-seconds are saved, they consume additional 100KB for every 2 minutes. The initial summaries will take about 3MB/hour, 70MB/day, 2GB/month, and 24GB/year but these detailed summaries can be discarded after a certain period.

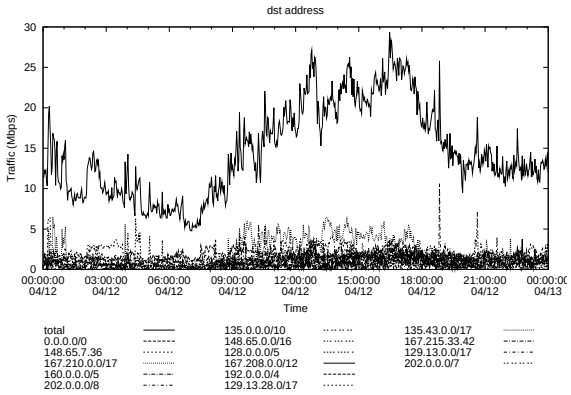

Figure 5. a graph plotting 1-day destination addresses

Plot Graph

Aguri supports a plot format output suitable to draw a plot graph. The plot format lists the counter values of the entries in a line; each line corresponds to a profiling period. It also supports conversion from byte-count to bits-per-second. A plot output is usually created from archived summaries and does not need to do in real-time. It is also needed to specify the number of entries in a plot. Thus, the plot generator uses a 2-phase algorithm which reads input files twice.

The first phase computes the cumulative byte count for each entry. At the end of the first phase, a sorted plot list is created, and the smallest entry is repeatedly aggregated until the number of nodes is reduced to the specified number. The second phase produces a plot format output for each period. For each period, if a node is not found in the plot list, it is aggregated to the nearest ancestor listed in the plot list. Hence, all counts are reflected to the plot.

Figure 5, 6 and 7 show examples of plot graphs taken from the trans-pacific link. The legend below the graph shows entries in the plot. Figure 5 plots destination addresses for 1 day on April 12, 2001, created from 2-minute summaries. Two individual addresses (148.65.7.36 and 167.215.33.42) are listed but there is no prominent address in terms of the bandwidth share.

Figure 6 plots source protocols for 10 days, from April 10th to 19th, 2001, created from 1-hour summaries. The graph captures daily fluctuations of the total traffic and the high ratio of HTTP. In Figure 6, there is a change in the daily traffic pattern on the 17th. By zooming into the 17th as shown in Figure 7, we can see unusual surges of ICMP. It is a *smurf* attack and this is the cause of the distortion in the daily traffic. We can identify the target address and the address range of the originators by looking into the corresponding address profiles. In fact, the corresponding source

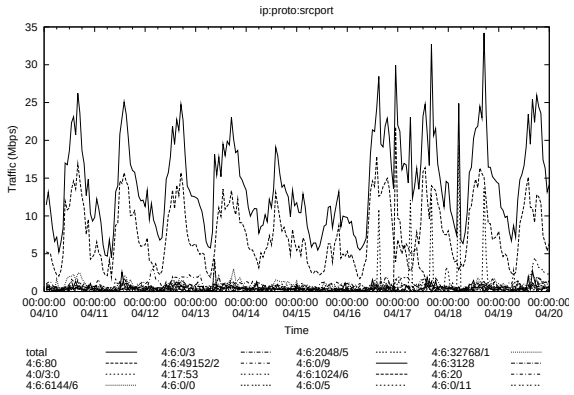

Figure 6. a graph plotting 10-day source protocols

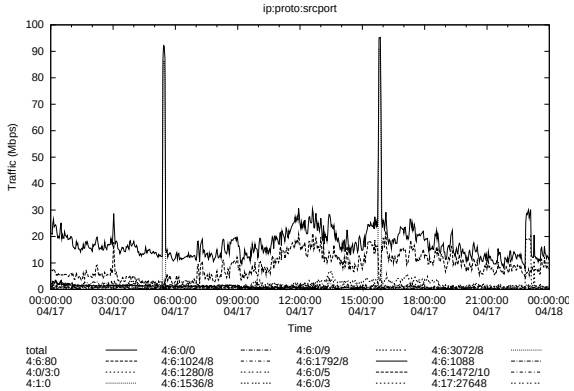

Figure 7. a graph zooming into April 17th

address profile has most traffic recorded from "0.0.0.0/1" and "208.0.0.0/6". This indicates the source addresses are widespread in the address space. On the other hand, the destination address profile shows that most traffic are targeted for an irc server within our network. This illustrates how plot graphs in different time scales can be used for trouble shooting.

Traffic Density Graph

Another graph format shows traffic density within the entire address space. From a summary, we can compute the traffic density in the address range of each aggregate, and create a time-series color graph. In a traffic density graph, the degree of traffic concentration is shown by colors and a change in traffic pattern is easily identified.

4 Related Work

MRTG [14] and its successor RRDtool [13] create time-series round-robin databases. They store numerical time-series data and automatically aggregate it into averages over time. Our idea of producing a summary from summaries is inspired by MRTG and RRDtool but differs in combining temporal aggregation with spatial aggregation.

Traditional flow-based monitoring tools such as NeTraMet [1] and FlowScan [15] require predefined rules to monitor a specific type of traffic. For example, in order to monitor HTTP traffic, they need to be instructed to identify TCP port 80. The approach with explicit and fixed rules has limitations on identifiable traffic types. Especially, it is a problem to cope with unknown protocols or DoS attacks.

Another approach is to report the top N flows by sorting the flow list [18, 3]. Although it does not need a rule set, there could be limitations on the maintainable number of flows or a flooding attack could easily overflow the list. Hence, it is not suitable for detecting DoS attacks. In our approach, a flooding attack may be able to reduce the resolution of the profile but the counter values are never lost. It is resilient to DoS attacks in addition to requiring no rules.

Dynamic identification of a flow is also addressed in the context of congestion control and DoS prevention. Floyd *et al.* in [8] argue on the need for end-to-end congestion control, and further, on the need for mechanisms in the network to detect and restrict unresponsive or high-bandwidth best-effort flows in times of congestion. They suggest to use the RED drop history as samples to identify misbehaving flows. The concept is known as a RED penalty-box [4].

This idea is further extended and detailed in order to cope with DDoS attacks and flash crowds [12]. It consists of a mechanism to identify aggregates, a local rate-limiter mechanism, and a pushback mechanism to propagate protective actions to neighbors. The proposed technique to identify high-bandwidth aggregates is based on the destination address in the drop history, and clusters the addresses into aggregates. The approach of identifying high-bandwidth aggregates and regulate them is similar to ours in the concept.

While their focus is to identify misbehaving flows, our focus is a traffic profiler which monitors and reports the network not only under congestion but all the time. Our observation is that a network point needing a protection mechanism is often a point to be monitored. Hence, it is practical to provide a combined solution both for performance and for simplicity. The combined method comes with visible monitoring outputs so that it could be advantageous to deployment.

5 Conclusion

We have described an aggregation technique for monitoring network traffic. We were in need of an adaptive traffic profiler to track long-term trend and to discover problems in our backbone network, and have developed a tool called aguri. Aguri adapts itself to spatial traffic distribution by aggregating small volume flows into aggregates, and achieves temporal aggregation by creating a summary of summaries applying the same algorithm to its outputs.

We have been using aguri for monitoring the WIDE backbone since February 2001, and found it useful for network operation. Especially, its ability to adapt to dynamic traffic and to visualize traffic in different time scales is powerful. If the traffic pattern becomes unusual, even if it is in an unpredictable way, it can be easily detected by looking at coarse grained plot graphs. Most of the information required for trouble shooting can be obtained by looking into finer grained summaries. If a DoS attack occurs, the target host as well as the type of the attack can be easily identified so that the operator can take prompt actions, for example, by setting filters at a border router.

There are a number of directions to improve the tool. We will continue to seek better visualization techniques. As the number of monitoring points increases, we need to automate trouble detection. Also, a management tool is needed for distributed monitoring in which a server collects and archives data from remote monitoring sites. Another interesting area is measurement-based active traffic control.

The implementation of aguri along with the related tools and other information is available from http://www.csl.sony.co.jp/~kjc/software.html.

References

[1] N. Brownlee. Traffic flow measurement: Experiences with NeTraMet. Request for Comments 2123, Internet Engineering Task Force, Mar. 1997.

[2] N. Brownlee, C. Mills, and G. Ruth. Traffic flow measurement: Architecture. Request for Comments 2722, Internet Engineering Task Force, Oct. 1999.

[3] K. Cho. Tele Traffic Tapper. http://www.csl.sony.co.jp/~kjc/software.html, 1996.

[4] K. Cho. Flow-valve: Embedding a safety-valve in red. In *Global Internet Symposium, Globecom*, pages 1753–1762, Dec. 1999.

[5] K. Cho, R. Kaizaki, and A. Kato. Aguri: An aggregation-based traffic profiler. In *QofIS2001*, pages 222–242, Coimbra, Portugal, Sept. 2001.

[6] K. Cho, K. Mitsuya, and A. Kato. Traffic data repository at the WIDE project. In *USENIX 2000 Annual Technical Conference: FREENIX Track*, pages 263–270, June 2000.

[7] K. C. Claffy, H.-W. Braun, and G. C. Polyzos. A parameterizable methodology for internet traffic flow profiling. *IEEE Journal of Selected Areas in Communications*, 13(8):1481–1494, 1995.

[8] S. Floyd and K. Fall. Promoting the use of end-to-end congestion control in the internet. *IEEE/ACM Transaction on Networking*, 7(4):458–472, Aug. 1999.

[9] V. Jacobson, C. Leres, and S. McCanne. tcpdump. ftp://ftp.ee.lbl.gov/, 1989.

[10] V. Jacobson, C. Leres, and S. McCanne. libpcap. ftp://ftp.ee.lbl.gov/, 1994.

[11] K. Keys, D. Moore, R. Koga, E. Lagache, M. Tesch, and K. Claffy. The architecture of the CoralReef internet traffic monitoring software suite. In *PAM 2001*, Amsterdam, The Netherlands, Apr. 2001.

[12] R. Mahajan, S. M. Bellovin, S. Floyd, J. Ioannidis, V. Paxson, and S. Shenker. Controlling high bandwidth aggregates in the network. *draft paper*, Feb. 2001.

[13] T. Oetiker. RRDtool: Round Robin Database Tool. http://ee-staff.ethz.ch/~oetiker/webtools/rrdtool/.

[14] T. Oetiker. MRTG: The multi router traffic grapher. In *USENIX LISA Conference*, pages 141–147, Boston, MA, Dec. 1998.

[15] D. Ponka. FlowScan: A network traffic flow reporting and visualization tool. In *USENIX LISA Conference*, New Orleans, LA, Dec. 2000.

[16] J. T. Robinson and M. V. Devarakonda. Data cache management using frequency-based replacement. In *SIGMETRICS Conference on Measurement and Modeling of Computer Systems*, pages 134–142, May 1990.

[17] K. Sklower. A tree-based packet routing table for berkeley UNIX. In *USENIX Winter Conference*, Dallas, Texas, Jan. 1991.

[18] S. Waldbusser. Remote network monitoring management information base. Request for Comments 1757, Internet Engineering Task Force, Feb. 1995.