


(19) **United States**(12) **Patent Application Publication** (10) **Pub. No.: US 2007/0076606 A1**
(43) **Pub. Date:** **Apr. 5, 2007**(54) **STATISTICAL TRACE-BASED METHODS
FOR REAL-TIME TRAFFIC
CLASSIFICATION**(75) Inventors: **Wladyslaw Olesinski**, Kanata (CA);
Peter Rabinovitch, Kanata (CA)Correspondence Address:
KRAMER & AMADO, P.C.
1725 DUKE STREET
SUITE 240
ALEXANDRIA, VA 22314 (US)(73) Assignee: **ALCATEL**, Paris (FR)(21) Appl. No.: **11/226,328**(22) Filed: **Sep. 15, 2005****Publication Classification**(51) **Int. Cl.**
H04L 12/26 (2006.01)(52) **U.S. Cl.** **370/230**(57) **ABSTRACT**

Apparatus and methods for real-time traffic classification based on off-line determined traffic classification rules are provided. Traces of real traffic are obtained and subjected to statistical analysis. The statistical analysis identifies domain the a multidimensional space of characteristic traffic parameters. Classification rules associated with the identified domains are derived and provided to traffic classification points for real-time traffic classification. Traffic classification points, typically edge network nodes, sample packets in aggregate streams with a predetermined probability. Statistical information regarding the sampled flows is tracked in a table, the number of time a flow was sampled providing a probabilistic measure of the flow's duration before the flow terminates. The table entries, which predominantly track high bandwidth flows, are subjected to the classification rules for real-time classification of the sampled flows. Optionally, rules include an action to be taken in respect of flows having characteristics matching thereof. Advantages are derived from low overhead on-line real-time classification of high-bandwidth flows at low overheads before flow termination.






FIG. 1

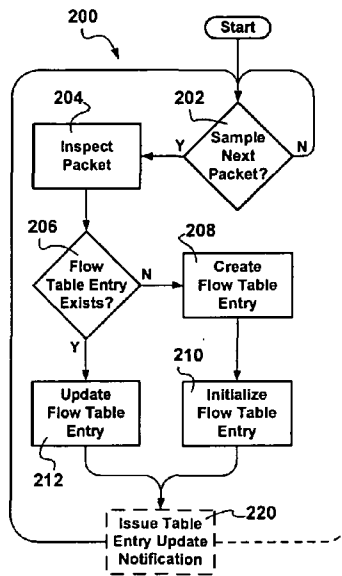


FIG. 2

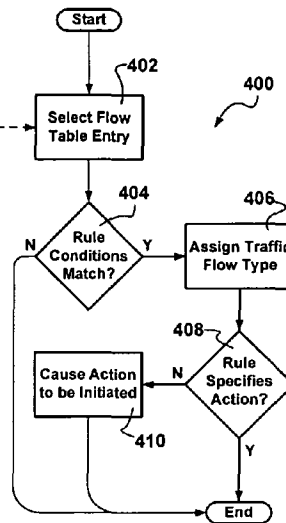


FIG. 4

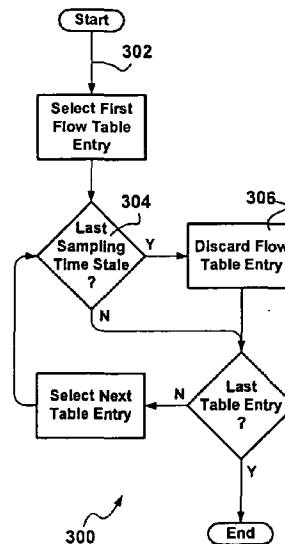


FIG. 3

STATISTICAL TRACE-BASED METHODS FOR REAL-TIME TRAFFIC CLASSIFICATION

FIELD OF THE INVENTION

[0001] The invention relates to content delivery at the edge of communications networks, and in particular methods and apparatus providing real-time trace-based traffic classification.

BACKGROUND OF THE INVENTION

[0002] Traffic classification is important for many reasons in delivering content to customers at the edge of communications networks. For example, Quality of Service (QoS) requires the traffic to be segregated first in order to assign packets to particular Classes of Service (CoS). A network operator can provide a different level of service to each class as well as a pricing structure.

[0003] Knowledge of traffic characteristics can help optimize the usage of the communications network infrastructure employed, and can help ensure a desired level of performance for applications/services important to the customers. The intention has always been that application requirements be considered in offering a level of service. Traditional methods of traffic detection and classification rely on monitoring logical port specifications typically carried in packet headers as, in the past, applications and/or services were, in a sense, assigned well known logical ports.

[0004] A large percentage of the traffic conveyed by communications networks today consists of peer-to-peer (P2P) traffic. Because peer-to-peer traffic is conveyed between pairs of customer network nodes, it is not necessary that a well known logical port be allocated, reserved, and assigned to traffic generated by applications generating peer-to-peer traffic and/or applications retrieving peer-to-peer content. Therefore known approaches to traffic classification are no longer valid as logical ports are undefined for peer-to-peer applications and/or logical ports may be dynamically allocated as needed such in the case of the standard File Transfer Protocol (FTP) and others.

[0005] Peer-to-peer content exchange techniques are increasingly being used to convey without permission content subject to intellectual property protection, such as music and movies. Network operators are under an increasing regulatory pressure to detect peer-to-peer traffic and to control illicit peer-to-peer traffic, while rogue users are seeking ways to defy traffic classification to avoid detection.

[0006] Besides peer-to-peer traffic detection, means and methods are being sought on a continual basis for detecting short duration traffic flows to help identify possible intrusions such as, but not limited to, Denial of Service (DOS) attacks.

[0007] Statistical billing is another domain in which knowledge of traffic characteristics is necessary. Network operators increasingly employ resource utilization measurements as a component in determining customer charges.

[0008] Returning to peer-to-peer traffic detection, not all peer-to-peer traffic is illicit: in view of the high levels of resource utilization demanded by peer-to-peer traffic, network operators may want to charge customers generating peer-to-peer traffic and retrieving peer-to-peer content more

for their high bandwidth usage. Resource utilization alone is not always an adequate traffic characteristic differentiator as in many instances content conveyed to, and received from, multiple customers is aggregated at the managed edge and within the managed transport communications network.

[0009] Attempts to characterize traffic, to detect traffic types, with a view of classifying traffic, include Deep Packet Inspection (DPI) techniques. Deep packet inspection techniques are described by Sen S., Spatscheck O. and Wang D. in "Accurate, Scalable In-Network Identification of P2P Traffic Using Application Signatures", Proceedings of the 13th international conference on World Wide Web, New York, NY, 2004; and by Karagiannis T., Broido A., Faloutsos M., Claffy K. in "Transport layer identification of P2P traffic", Proceedings of the 4th ACM SIGCOMM Conference on Internet Measurement, Taormina, Sicily, Italy, 2004.

[0010] Proposed deep packet inspection techniques, as the name suggests, assume the availability of unlimited resources to inspect entire packets to the perform packet characterization. Therefore deep packet inspection incurs high processing overheads and is subject to high costs. Deep packet inspection also suffers from a complexity associated with the requirement of inspecting packet payloads at high line rates. For certainty, deep packet inspection is not suited at all for typical high throughput communications network nodes deployed in current communications networks. Deep packet inspection also suffers from a high maintenance overhead as the detection techniques rely on signatures, peer-to-peer applications, especially, are known for concealing their identities—a deep packet inspection detection signature that provides conclusive detection now may not work in the future, and another conclusive signature would have to be found and coded therein.

[0011] Traffic classification means and methods are being actively sought by network operators in order to determine the types of traffic present in a managed communications network for traffic and network engineering purposes, on-line marking of packets, quality of service assessment/assurance, billing, etc. In view of impending regulatory pressures, efficient detection and classification of peer-to-peer traffic is especially desired, as peer-to-peer traffic consumes large, disproportional percentages of bandwidth and other communication network resources. Network operators have to employ a combination of: peer-to-peer traffic control in order to reserve network resources for other types of traffic, charge peer-to-peer users different rates to curb behavior, and/or even block peer-to-peer completely in accordance with regulations imposed on network operators. There therefore is a need to solve the above mentioned issues to provide traffic classification means and methods which avoid the complexities of deep packet inspection and the pitfalls of logical port based packet classification.

SUMMARY OF THE INVENTION

[0012] In accordance with an aspect of the invention, a packet flow classification apparatus for on-line real-time traffic flow classification at a communications network node is provided. Packet sampling means randomly selects packets from an aggregate flow with a pre-determined sampling probability. Packet inspection means determine the packet size of each sampled packet and obtains the flow identification of the sampled traffic flow with which the sampled

packet is associated. A sampled flow information table has sampled flow table entries for storing real-time sampled flow statistical information. Flow information tracking means maintain the flow information table in real-time. And, a packet classifier classifies sampled traffic flows on-line in real-time based on a group of classification rules trained off-line on statistical trace traffic flow information.

[0013] In accordance with another aspect of the invention, a packet flow classification system for on-line real-time traffic flow classification of a plurality of traffic flows conveyed through a network node of a communications network is provided. The packet flow classification system includes: a group of classification rules trained off-line on statistical trace traffic flow information, at least one traffic flow monitor at the communications network node, and a packet classifier for classifying sampled traffic flows on-line in real-time based on the group of off-line trained classification rules. The traffic flow monitor includes: packet sampling means for randomly selecting packets from an aggregate flow with a pre-determined sampling probability, packet inspection means for determining the packet size of each sampled packet and for obtaining flow identification of the sampled traffic flow with which the sampled packet is associated, a sampled flow information table having sampled flow table entries for storing real-time sampled flow statistical information, and flow information tracking means for maintaining the flow information table in real-time.

[0014] In accordance with yet another aspect of the invention, a method of classifying traffic flows on-line in real-time based at least one classification rule trained off-line on statistical traffic flow information is provided. Packets are randomly sampled from an aggregate flow with a predetermined sampling probability. The packet size of each sampled packet is extracted. A flow identifier of the sampled traffic flow with which the sampled packet is associated is obtained. Information regarding the sampled flow is tracked in a sampled flow information table entry. And, the sampled traffic flow is classified in real-time by subjecting the information tracked in the flow table entry to the at least one classification rule.

[0015] Advantages are derived from simple, low overhead and inexpensive on-line real-time classification of high-bandwidth flows at low overheads before flow termination.

BRIEF DESCRIPTION OF THE DRAWINGS

[0016] The features and advantages of the invention will become more apparent from the following detailed description of the exemplary embodiment(s) with reference to the attached diagrams wherein:

[0017] FIG. 1 is a schematic diagram showing an edge communications network element sampling traffic sporadically and tracking statistical flow information in a table in accordance with the exemplary embodiment of the invention;

[0018] FIG. 2 is a flow diagram showing process steps in sampling traffic sporadically and tracking statistical flow information in a table in accordance with the exemplary embodiment of the invention;

[0019] FIG. 3 is a flow diagram showing process steps of a clean-up process ensuring that the statistical information in

a flow information table is current, in accordance with the exemplary embodiment of the invention; and

[0020] FIG. 4 is a flow diagram showing process steps of an on-line real-time classification process classifying high-bandwidth sampled flows in accordance with the exemplary embodiment of the invention.

[0021] It will be noted that in the attached diagrams like features bear similar labels.

DETAILED DESCRIPTION OF THE EMBODIMENTS

[0022] The realities of Internet service provisioning to customers are such that service level agreements are described in terms of expected aggregate traffic characteristics with the assumption that most of the user traffic is highly bursty and relatively low bandwidth such as the occasional email, intermittent web page download followed by a reading period, and the infrequent electronic bank transaction. Although the equipment is prevalent, video conferencing is relatively rare. Service level agreements include enough long-term transport bandwidth for comparatively higher bandwidth netradio audio streaming. Customers are assumed to be nice and occasional transgressions rarely translate into higher bills at the end of the month. It is assumed that nice customers do not listen to netradio, nor download MP3's from traceable and reputable sources, 24/7. At the same time, in view of the intense competition in communications, the available transport bandwidth in the core of the managed communications network is oversold.

[0023] Problems arise when customers engage in illicit/regulated activities such as exchanging large amounts of content subject to intellectual property protection. Such rogue customers do not want to pay higher fees for levels of service which would provide them with increased bandwidth in order not to attract attention to themselves. Sophisticated rogue customers are willing to put up with sending and/or receiving content at transfer rates well below aggregate service level agreement limits for long periods of time. Network operators are faced with the conundrum that: rogue customers do not violate their service level agreements; the fact that the rogue customers have not signed up at a higher level of service the resources are overused; and because of the bandwidth is oversold, services provisioned to nice customers are being impacted. Therefore given that input traffic from multiple customers and traffic output to multiple customers is typically aggregated on edge, aggregate traffic metrics only point out that network resources are utilized to a very high degree and that the average and the typical customer is nice.

[0024] With the prevalence of peer-to-peer activity, it does not make business sense for network operators to deny services to any user engaging in peer-to-peer file sharing, after all some peer-to-peer traffic such as a video conference is legitimate. What is desired is to identify regulated peer-to-peer traffic in order to reduce the allocation of network resources thereto thereby making the network unfriendly only to questionable/undesirable peer-to-peer traffic.

[0025] Current traffic characterization techniques, in order to be effective, employ a determination of particular traffic flows and their duration to determine the amount of resources expended. Because of the willingness of rogue

users to make peer-to-peer traffic compliant over the short term with service level agreements subscribed to at the cost of long duration uploads/downloads, knowing the duration of a traffic flow has been found to be very important in characterizing traffic. The trouble is that using current techniques, the duration of a traffic flow can only be found after the traffic flow ends, when is too late to effect any control over the traffic flow. For this very reason, current deep packet inspection techniques, which are designed to detect the initiation and the termination of a traffic flow, are inadequate as a trigger to real-time peer-to-peer traffic control as the traffic classification is provided after the termination of the monitored activity.

[0026] Having identified the problem, real-time traffic identification is needed to provide a measure of how the available network resources are utilized by, and partitioned to, rogue customers in real-time. From a business point of view, any solution has to assume that customers are nice. This is necessary both from the customer relations point of view and because assuming the converse, would require a prohibitive amount of resources to be devoted to traffic monitoring.

[0027] In "Class-of-Service Mapping for QoS: A statistical signature-based approach to IP traffic classification", ACM SIGCOMM Internet Measurement Workshop, Taormina, Sicily, Italy, 2004, which is incorporated herein by reference, Roughan M., Subhabrata S., Spatscheck O. and Duffield N., address the fact that although various mechanisms exist for providing QoS, QoS has yet to be widely deployed. Roughan et al. are of the opinion that employing previously known techniques to map traffic to QoS classes is prohibitive due to the high overheads incurred. Falling short of providing real-time traffic classification, Roughan et al. predicate their off-line traffic classification solution on the fact that it would be unrealistic for effective traffic classification to inspect every packet, and propose an off-line trace-based statistical method of traffic characterization. Roughan et al. confirm that flow durations are essential in characterizing and distinguishing between different types of traffic. However, just like deep packet inspection techniques, the trace-based methods proposed by Roughan et al. only obtain the flow duration after each flow terminates and are therefore inadequate for real-time traffic characterization.

[0028] Other relevant research in the art of packet queuing includes an proposal by Psounis K., Gosh A. and Prabhakar B. entitled "SIFT: A low-complexity scheduler for reducing flow delays in the Internet", Technical report, CENG-2004-01, USC, which is incorporated herein by reference, and describes an algorithm for identifying high bandwidth flows in order to queue packets of identified high bandwidth flows in a special queue. The SIFT proposal includes sampling packets with a pre-selected low probability. With the assumption that most low bandwidth flows consist of few packets, very few low bandwidth flows would be sampled. Conversely high bandwidth flows would be sampled with greater certainty. The flow identifier of each sampled flow is provided to a packet classifier/queue manager which queues every subsequent packet bearing one of the provided flow identifiers into the special queue. The proposed use of the special queue(s) would be prohibitive in terms of the necessary resources for high throughput deployments.

[0029] In accordance with the exemplary embodiment of the invention, information is gathered about real typical traffic patterns for statistically relevant periods of time. The outcome of information gathering step is a trace of traffic (packet headers and, possibly, portions of payloads) and perhaps a collection of gathered and/or derived statistics. The trace traffic information is gathered with the intent of subjecting thereof to off-line traffic characterization training.

[0030] Because the proposed traffic characterization is performed off-line, information about known and/or determinable traffic flow types may also be used as inputs to a rule creation process referred to as off-line training. Using the trace information, different applications generating/consuming the traffic are identified. Diverse off-line methods can be used for this purpose, for example, the methods described in the above referenced prior art Sen and Karagiannis and other deep packet inspection methods to the extent that portions of payloads have been obtained. This step associates traffic flows with the applications.

[0031] Classes of traffic, or traffic types, are defined. For example, Roughan et al. propose the following exemplary traffic classes:

[0032] interactive, including: remote login sessions, interactive Web content access, etc;

[0033] bulk data transfer, including: FTP, music downloads, peer-to-peer traffic, etc.;

[0034] streaming, including video conferencing, netradio, etc.; and

[0035] transactional, including: distributed database access.

[0036] Traffic characterization means are employed off-line to characterize the traffic flows based on the gathered trace information and the associated statistics, to create off-line rules based on statistical properties of the gathered information and the application associativity. The outcome of the characterization step is a relation between traffic statistics (packet sizes, flow durations, etc.), application associativity and the traffic classes the traffic flows belong to. In creating the rules, Roughan et al. propose subjecting the traffic flow statistical parameters to statistical analysis in order to define domains in the multi-dimensional space of statistical flow parameters. Subjecting traffic flow statistical parameters to statistical analysis is referred herein as training. The domains correspond to statistically distinct traffic classes. Statistical analysis methods include, but are not limited to, Nearest Neighbour (NN), Linear Discriminant Analysis (LDA), etc. Given traffic classes, gathered statistics, trace information, and application associativity, a set of rules are trained for classifying future data is determined.

[0037] For example assuming that only two statistics: average packet sizes and session durations, are all that is needed to characterize traffic, a rule can indicate that if an average duration of a traffic flow is greater than 40 seconds, and average packets of that traffic flow do not exceed 300 bytes, then the traffic flow most likely belongs to interactive class (e.g., a Telnet session). As another example, peer-to-peer flows, based on the number of packets conveyed, statistically fall into bulk data transfer flows and streaming flows, however average packet sizes would characterize peer-to-peer flows as bulk data transfer flows regardless of

the application or logical port used to convey the content. Beside flow durations, traffic flows can be characterized based on statistical traffic flow parameters such as: average/median packet size, packet size variance, root-means-square packet size, largest packet sampled so far, shortest packet sampled so far, average/median inter-packet arrival delay, inter-packet arrival delay variance, bytes per flow, packets per flow, etc.

[0038] In accordance with an exemplary implementation of the exemplary embodiment of the invention, classification rules pertaining to expected/uninteresting traffic patterns may be deleted from the off-line determined set of classification rules. Reducing the number of classification rules provides desirable overhead reductions.

[0039] Proposed real-time methods include real-time statistics collection and real-time traffic classification.

[0040] In accordance with the exemplary embodiment of the invention, random packet sampling techniques are employed on a monitored aggregated flow irrespective of the individual aggregated flows contained therein for the purpose of estimating flow durations of individual sampled flows. Random packet sampling techniques conform to the desired assumption that initially all customers' traffic flows are nice.

[0041] In accordance with the exemplary embodiment of the invention, network elements shown in FIG. 1 such as, but not limited to, edge switching equipment and routing equipment, include packet selection means 150, packet inspection means 152, statistical information tracking means 154, and a flow information table 100, having sampled flow table entries 102. Each sampled flow entry 102 exemplary includes fields for the flow identifier 104, the number of packets sampled 106, the cumulative amount of content sampled 108, sampling time of the last packet 110. All information necessary to populate and update the fields 104, 106, 108 and 110 can either be extracted by the packet inspection means 152 or derived by the statistical information tracking means 154 from information extracted from sampled packet headers. For certainty, dependent on the implementation, the packet selection means 150, the packet inspection means 152, the statistical information tracking means 154, and the flow information table 100, without limiting the invention, can be associated with a physical port, a logical port, a group of physical/logical ports, all ports, etc. Also, depending on the implementation, the proposed traffic monitoring may only be performed on best-effort traffic and/or available bit rate traffic conveyed by the implementing equipment.

[0042] In accordance with the exemplary embodiment of the invention, a real-time traffic monitoring process 200 shown in FIG. 2 maintains the flow information table 100. Based on a preset sampling probability, the traffic monitoring process 200 determines 202 whether to sample the next packet upon arrival. If the next packet is to be sampled, then the process 200, waits for the next packet to arrive, and inspects 204 the received packet to extract information such as, but not limited to, flow identification, and packet size. Typically, the packet inspection 204 is typically limited to reading at most the packet header, application layer connection information hidden in packet payloads is typically not searched for. Existing packet inspection means otherwise used in packet processing may be reused.

[0043] If a flow table entry 102 for the identified flow does not exist 206 in the flow information table 100, then a flow table entry 102 is created 208 and the flow table entry is initialized 210. Initializing 210 the flow table entry 102 includes, but is not limited to, filing in the flow identifier field 104, setting the number of sampled packets 106 to 1, setting the cumulative amount of content sampled 108 to the size of the sampled packet, and writing the sampling time in field 110.

[0044] If a flow table entry 102 exists for the identified flow, then the fields of the table entry 102 are updated 212 by: incrementing the number of sampled packets 106 by 1, adding the size of the sampled packet to the cumulative amount of content sampled 108, and overwriting field 110 with the sampling time.

[0045] Having created 208 or updated 212 the flow table entry 102, the monitoring process 200 resumes from step 202. The resumption may include the issuance 220 of a notification that the particular flow entry 102 has been updated.

[0046] In accordance with the exemplary embodiment of the invention, the number of times a traffic flow is sampled, regardless of the number of actual packets associated with the corresponding traffic flow being conveyed, is representative of the duration of the traffic flow so far. In view of the off-line training, the number of packets sampled so far, perhaps together with the other statistical values tracked in the flow table entry 102, is representative of the actual duration of the traffic flow. Therefore duration of a sampled flow can be forecasted before the traffic flow terminates within statistical certainty when the field values in the corresponding flow entry table match a rule.

[0047] Tracking the cumulative amount of sampled content and the number of sampled packets is equivalent to tracking the average packet size without requiring performing division as packets are sampled in real-time at substantial overhead reductions.

[0048] An exemplary parallel flow information table clean-up process 300, shown in FIG. 3 inspects 304 sampling times specified in flow table entries 102 and discards 306 stale entries 102. The clean-up process 300 executes 302 in accordance with a clean-up discipline, for example periodically.

[0049] The monitoring process 200 and the clean-up process 300, identify the existence of, and monitor, suspect traffic flows on-line and in real-time. The number of sampled packets of each flow entry 102 in the flow table 100 provides an estimation of the duration of each flow. And, the clean-up process 300 makes up for the lack of a determination of the exact time when monitored flows terminate. Having real-time statistical information regarding current high bandwidth flows presents the network operator with the most plausible flows to consider in identifying illicit/regulated content transfers regardless of the rogue customers' attempts to foil detection.

[0050] In accordance with an exemplary implementation of the exemplary embodiment of the invention, an on-line real-time traffic classification entity 250, typically associated with the traffic control point 252, is provided with the traffic classification rules and performs an on-line real-time classification process 400 as exemplary shown in FIG. 4.

Accordingly, on-line real-time traffic classification is achieved by subjecting **404** the suspect traffic flows monitored via the flow information table **100** to the off-line trained classification rules. Simply put, subjecting flow entry **102** to each rule answers the question whether the sampled flow, to which the flow entry **102** corresponds, has statistical characteristics which would locate the traffic flow in the multi-dimensional domain expressed in the rule. Without limiting the invention, the on-line real-time classification process **400**, may be triggered **402** via notifications **220** or in accordance with a discipline, for example on a schedule or periodically. If there is a hit **404** in respect of a traffic flow, the classification rule hit specifies **406** the traffic class/type. Accordingly, information about the current state of the monitored communications network is provided allow the network operator to use this information to take further action in respect of specific flows in real-time if and when necessary.

[0051] In accordance with the exemplary embodiment of the invention, the off-line trained traffic classification rules are also used to identify illicit/regulated content transfers from the suspect flows identified. Depending on implementation, and without limiting the invention, rule hits **404** may either be (**410**) logged, alarms may be raised, and/or traffic control may be enforced at a traffic control point **252**. Depending on implementation, and/or if the rule applies to illicit/regulated traffic, the rules themselves may also specify **408** appropriate actions.

[0052] In accordance with an exemplary implementation of the invention, actions specified **408** via the classification rules, without limiting the invention, include, marking packets of the sampled flows, assignment of packet processing priorities, assignment of sampled flow packets to specific classes-of-service, special processing of different traffic types, precise monitoring of traffic, etc.; wherein special packet processing includes discarding packets of the sampled flow.

[0053] In accordance with exemplary implementations of the exemplary embodiment of the invention, depending on the content transport technology deployed, traffic control may be effected on the actual edge network node as in the exemplary case of Internet Protocol deployments, or effected centrally as in the exemplary case of A synchronous Transfer Mode deployments.

[0054] If the traffic classification entity **250** is remote with respect to the network node maintaining the flow information table **100**, typically co-located with other centralized network provisioning functions, the statistical information in the flow information table **100** is exported thereto in accordance with an export discipline and/or in response to a request from a network controller **254**. Because only high-bandwidth flows are sampled, and because the high-bandwidth flows are sampled sparingly, it is envisioned that the remote traffic classification entity **250** can be provided with appropriate notifications (**220**) for each change to the flow information table **100**. However, for high density deployments, multiple flow information table change notifications **220** may be accumulated and the accumulated notifications are transmitted at reduced management overheads. In an exemplary ATM deployment, statistical connection (flow) information is exported to an off-board network controller

254, the traffic classification entity **250** is associated with the network controller **254** which implements functions of the traffic control point **252**.

[0055] In accordance with an exemplary implementation of the exemplary embodiment of the invention, the traffic classification entity **250** is located within the network element sampling traffic, as is the traffic control point **252**. While such an arrangement would introduce a small and sporadic rule distribution overhead, it would reduce the larger statistical flow information reporting overhead necessary for centralized traffic classification described above. In accordance with another exemplary implementation of the exemplary embodiment of the invention, the network element sampling traffic has means for retrieving the group of off-line trained classification rules. In an exemplary IP deployment, the traffic classification entity **250** is implemented on a control card of a modular network node, existing functions of the control card representing the traffic control point **252**.

[0056] It was mentioned above that actions **408** specified via the classification rules include precise monitoring of traffic. It is envisioned that precise monitoring of traffic include deep packet inspection. For certainty, deep packet inspection, when necessary, would be performed after real-time traffic classification only on suspect traffic flows sampled, at substantial packet processing overhead reductions.

[0057] Accordingly, acquired knowledge derived from the off-line statistical analysis is used via the trained classification rules to perform on-line real-time classification of traffic in a managed communications network. For clarity, while deep packet inspection may be used off-line to identify traffic characteristics prior to training the classification rules, deep packet inspection is not used in performing on-line real-time traffic classification making the proposed solution suitable for high throughput applications.

[0058] Because the solution does not need to track customers, just flows irrespective of customers, the assessment preserves the customers' anonymity while any flow control exercised ensures that the network resources are shared by all customers in a fair fashion.

[0059] The embodiments presented are exemplary only and persons skilled in the art would appreciate that variations to the above described embodiments may be made without departing from the spirit of the invention. The scope of the invention is solely defined by the appended claims.

We claim:

1. A packet flow classification apparatus for on-line real-time traffic flow classification at a communications network node, the packet flow classification apparatus comprising:

- a. packet sampling means for randomly selecting packets from an aggregate flow with a pre-determined sampling probability;
- b. packet inspection means for determining the packet size of each sampled packet and for obtaining flow identification of the sampled traffic flow with which the sampled packet is associated;
- c. a sampled flow information table having sampled flow table entries for storing real-time sampled flow statistical information;

- d. flow information tracking means for maintaining the flow information table in real-time;
 - e. a group of classification rules trained off-line on statistical trace traffic flow information; and
 - f. a packet classifier for classifying sampled traffic flows on-line in real-time based on the group of off-line trained classification rules.
2. A packet flow classification apparatus as claimed in claim 1, wherein each flow table entry further stores the number of times the corresponding traffic flow was sampled; given the predetermined sampling probability, the number of times the traffic flow was sampled providing an indication of the duration of the traffic flow within a statistical certainty.
3. A packet flow classification apparatus as claimed in claim 1, wherein each flow table entry further stores one of an average packet size, a median packet size, a packet size variance, a root-means-square packet size, a size of the largest packet sampled so far, a size of the shortest packet sample so far, an average inter-packet arrival delay, a median inter-packet arrival delay, the inter-packet arrival delay variance, a number of sampled bytes, and a cumulative amount of traffic content sampled so far.
4. A packet flow classification apparatus as claimed in claim 1, wherein each flow table entry further stores the time of the last packet sampling.
5. A packet flow classification apparatus as claimed in claim 4, wherein the flow information table tracking means further includes means for discarding a flow table entry having a stale last packet sampling time.
6. A packet flow classification apparatus as claimed in claim 1, further comprising means for obtaining the group of off-line trained classification rules.
7. A packet flow classification apparatus as claimed in claim 1, further comprising notification means for notifying the packet classifier of one of the creation, an update, and the deletion of a flow table entry performed by the flow information tracking means.
8. A packet flow classification apparatus as claimed in claim 1, wherein the packet sampling means is associated with one of a port of the communications network node, a group of ports of the communications network node, a trunk, best-effort traffic processed at the communications network node, available bit rate traffic processed at the communications network node, the communications network node, and a control card of the communications network node.
9. A communications network node comprising the packet flow classification apparatus claimed in claim 1, wherein each classification rule further specifies an action to be taken on subsequent packets constituent of the classified sampled traffic flow, the communications network node further comprising traffic control means for implementing the action.
10. A communications network node comprising the packet flow classification apparatus claimed in claim 9, wherein the action comprises one of: marking, assigning a class-of-service to, ascribing a packet processing priority to, queuing, diverting, inspecting, and discarding subsequent packets associated with the classified sampled flow.
11. A packet flow classification system for on-line real-time traffic flow classification of a plurality of traffic flows conveyed through a network node of a communications network, the packet flow classification system comprising:
- a. a group of classification rules trained off-line on statistical trace traffic flow information;
 - b. at least one traffic flow monitor at the communications network node, the traffic flow monitor including:
 - i. packet sampling means for randomly selecting packets from an aggregate flow with a pre-determined sampling probability;
 - ii. packet inspection means for determining the packet size of each sampled packet and for obtaining flow identification of the sampled traffic flow with which the sampled packet is associated;
 - iii. a sampled flow information table having sampled flow table entries for storing real-time sampled flow statistical information; and
 - iv. flow information tracking means for maintaining the flow information table in real-time; and
 - c. a packet classifier for classifying sampled traffic flows on-line in real-time based on the group of off-line trained classification rules.
12. A packet flow classification system as claimed in claim 11, wherein each flow table entry further stores the number of times the corresponding traffic flow was sampled; given the predetermined sampling probability, the number of times the traffic flow was sampled providing an indication of the duration of the traffic flow within a statistical certainty.
13. A packet flow classification system as claimed in claim 11, wherein each flow table entry further stores one of an average packet size, a median packet size, a packet size variance, a root-means-square packet size, a size of the largest packet sampled so far, a size of the shortest packet sample so far, an average inter-packet arrival delay, a median inter-packet arrival delay, the inter-packet arrival delay variance, a number of sampled bytes, and a cumulative amount of traffic content sampled so far.
14. A packet flow classification system as claimed in claim 11, wherein each flow table entry further stores the time of the last packet sampling.
15. A packet flow classification system as claimed in claim 14, wherein the flow information table tracking means further includes means for discarding a flow table entry having a stale last packet sampling time.
16. A packet flow classification system as claimed in claim 11, wherein the packet classifier is associated with the communications network node, the packet flow classification system further comprising means for providing the communications network node with the group of classification rules.
17. A packet flow classification system as claimed in claim 11, wherein the packet classifier is associated with a network controller, the communications network node further comprising means for providing sampled traffic flow information stored in the sampled flow information table to one of the network controller and the classifier.
18. A packet flow classification system as claimed in claim 17, wherein the means for providing traffic flow information further comprises notification means for notifying one of the network controller and the classifier of one of flow table entry creation, update, and deletion performed by the flow information tracking means.
19. A packet flow classification system as claimed in claim 18, wherein one of the network controller and the classifier further comprises retrieval means for retrieving information stored in the traffic sampled flow information table.

20. A packet flow classification system as claimed in claim 11, each classification rule further specifying an action to be taken on subsequent packets constituent of the classified sampled traffic flow, the packet flow classification system further comprising traffic control means associated with one of the network controller and the communications network node for implementing the action.

21. A packet flow classification system as claimed in claim 20, wherein the action comprises one of: marking, assigning a class-of-service to, ascribing a packet processing priority to, queuing, diverting, inspecting, and discarding subsequent packets associated with the classified sampled flow.

22. A packet flow classification system as claimed in claim 11, wherein the packet sampling means is associated with one of a port of the communications network node, a group of ports of the communications network node, a trunk, best-effort traffic processed at the communications network node, available bit rate traffic processed at the communications network node, the communications network node, and a control card of the communications network node.

23. A method of classifying packet traffic flows on-line in real-time based at least one classification rule trained off-line on statistical traffic flow information, the method comprising the steps of:

- a. randomly sampling packets, with a predetermined sampling probability, from an aggregate of a plurality of flows conveyed via a communications network node;
- b. extracting the packet size of each sampled packet;
- c. obtaining a flow identifier of the sampled traffic flow with which the sampled packet is associated;
- d. tracking information regarding the sampled flow in a sampled flow information table entry of a sampled flow table held at the communications network node; and
- e. classifying the sampled traffic flow in real-time by subjecting the information tracked in the flow table entry to the at least one classification rule.

24. A method as claimed in claim 23, wherein each flow table entry stores the number of times the corresponding flow was sampled, tracking information in the flow information table entry further comprises increasing the number of times the flow was sampled, the number of times the corresponding flow was sampled being indicative of the flow duration within a statistical certainty.

25. A method as claimed in claim 23, wherein tracking information in the flow information table entry further comprises storing in the flow table entry one of an average packet size, a median packet size, a packet size variance, a root-means-square packet size, a size of the largest packet

sampled so far, a size of the shortest packet sample so far, an average inter-packet arrival delay, a median inter-packet arrival delay, the inter-packet arrival delay variance, a number of sampled bytes, and a cumulative amount of traffic content sampled so far.

26. A method as claimed in claim 23, wherein each sampled flow table entry further stores the time of the last packet sampling, the method further comprising discarding a sampled flow table entry having a stale last packet sampling time.

27. A method as claimed in claim 23, wherein subjecting the tracked information in the flow table entry to the classification rule, the method further comprises determining whether the information tracked in the flow table entry specifies that the corresponding sampled traffic flow is one having statistical characteristics locating the sampled traffic flow in the domain defined by at least one statistical parameter as specified via the classification rule.

28. A method as claimed in claim 23, wherein distributed flow classification means are employed in classifying sampled traffic flows, the method further comprising obtaining the at least one off-line trained classification rule from a location in the communications network.

29. A method as claimed in claim 23, wherein tracking sampled flow information further includes one of creating, updating and deleting the flow table entry.

30. A method as claimed in claim 29, further comprising issuing a notification specifying one of the creation, the update, and the deletion of a flow table entry, the classifying step being responsive to the notification.

31. A method as claimed in claim 30, wherein centralized flow classification means are employed in classifying sampled traffic flows, the method further comprising providing the centralized flow classification means with sampled traffic flow information held in a sampled flow table entry.

32. A method as claimed in claim 31, further comprising requesting sampled traffic flow information held in the sampled flow table entry specified in the notification.

33. A method as claimed in claim 23, wherein each classification rule further specifies an action to be taken on subsequent packets constituent of the classified sampled traffic flow, the method further comprising implementing the action.

34. A method as claimed in claim 33, wherein implementing the action further comprises one of: marking, assigning a class-of-service to, ascribing a packet processing priority to, queuing, diverting, inspecting, and discarding subsequent packets associated with the classified sampled flow.

* * * * *