Source characteristics for traffic classification in Differentiated Services type of networks

Marko Luoma, Mika Ilvesmäki and Markus Peuhkuri

Helsinki University of Technology,
P.O. Box 3000,
02015 HUT, Finland

ABSTRACT

Differentiated services (DS) is a mechanism by which network service providers can offer differing levels of network service to different traffic, and in doing so they provide quality of service (QoS) to their customers. The advantage of DS is that many traffic flows can be aggregated to a small number of classes, thereby simplifying the processing and storage associated with packet classification and conditioning. In addition, there is no signaling state or related processing required in the DS network since QoS is invoked on a packet-by-packet basis. To succeed in differentiation, router must somehow decide which packet should go into which class. This decision can be based on inherent characteristics of traffic source or application or some other networking perspective. Traffic in Internet can be broadly categorized into two different classes: interactive and bulk. Interactive traffic is seen to be formed from conversational applications needing continuos attention from the user like IP telephony, video conferencing, collaborative application (shared white board etc) and applications forming a group of which some have interactive nature as an application (web, remote shell) or require some level consistency in delivery time like streaming audio and video. Bulk category takes all what is left from previous ones. These left overs are usually related to background processes like file transfer (FTP) or messaging (email).

In this paper we will present a survey of ways to do quality differentiation and packet classification. A group of classification mechanisms are investigated more closely, namely measurement based classifiers.

Keywords: Differentiated Services, Integrated Services, Quality differentation, Traffic Classification

1. INTRODUCTION

In multiservice networking where the application base is diversified, the notion of quality becomes important. Quality is usually used to describe the process of delivering data to a selected set of applications in a reliable manner within certain parametric boundaries (QoS) or in a manner which is somewhat better than what other applications or users are receiving (CoS).

New Internet applications are born in an ever increasing pace to help people communicate in the Internet using voice and video. The new applications, with continuous media playback properties, create new kind of traffic which requires better service levels from the network than can be delivered today. In the contemporary Internet there exists no quality differentiation; not even at the level of separating traffic based on the transport protocol. Lack of differentiation is becoming a problem for Internet users - the transport protocols (TCP and UDP) interfere with each other when sharing a single buffer in the IP router and thus reduce the quality of communication they are offering.

There have to be mechanisms for differentiating traffic for the network to be able to offer different levels of service to applications. These mechanisms, generally referred to as traffic classifiers, operate on packet, flow and application level. In this paper we present traffic classifier mechanisms based on the transport level identifiers targeted to networks which aim to offer homogenous classification of traffic.

Further author information: (Send correspondence to M.L)
M.L: E-mail: marko.luoma@hut.fi
M.I: E-mail: lynx@tct.hut.fi
M.P: E-mail: puhuri@cc.hut.fi

2. INTERNET SERVICE

The current Internet offers a single class of best-effort service; that is, there is no admission control and the network offers no assurance about when, or even if, packets will be delivered. Many of the real-time applications do not work adequately in this best effort environment. These applications, often implemented on top of UDP, interfere with data applications as they compete of the common buffer and link capacities. These interferences are born from the fact that UDP, as it cannot control the sending rate of the source in the presence of congestion, cannot back off during the times of congestion.

To offer classes of services with different quality targets, two service architectures have been designed for the future Internet: Integrated Services and Differentiated Services. These service architectures approach the problem from different viewpoints: Differentiated Services from the minimal, aggregated guarantee approach and the Integrated Services from the concept of per flow quality quarantee in the network.

2.1. Integrated Services

Integrated Services (IS or IntServ)[1–3] is a service architecture where the user initiates the requests for the service level he needs. In practice, the application explicitly request the service level as it prepares to send data. In other words, the network offers a set of service classes and the application indicates to the network which service class it wants and what are the related resource reservation parameters. The mapping of the application and the service class is thus under the control of the application and the user.

In Integrated Services, network keeps track of the state of each individual connection through the network. This tracking makes the process of state reservation for the resources work through a proper control procedure. The reservations are done periodically in order to refresh the reservation state. Refresh requests are not processed unless they contain changes to the original reservation, thus making the state soft.

In Integrated Services, where service class and quality is under control of user, the network must provide some system to encourage users to request a proper service class for their application. Pricing of the network services is one approach. Charging more for the higher quality service will ensure that only the most performance-sensitive applications will request higher service levels.

2.2. Differentiated Services

Differentiated Services (DS or DiffServ)[4] is a service architecture where either the user explicitly or the network implicitly chooses the appropriate service class for the flow without actually reserving resources to individual flows. In this case the application sends its packets, possibly without stating anything about its service requirements, and the network then classifies the packets into the proper service class and handles them then accordingly. Differentiated Services based network is thus stateless, much like the current Internet is. Stateless refers to the fact that there is no knowledge of individual connections inside the network, only the aggregate behavior of traffic flows is examined.

One of the advantages of this approach is that there is no service level negotiation in the network. This way applications do not need to specify their desired service level to the network, and the network does not, in turn, need to describe the delivered service to the application. Also, since there is no explicit commitment to a given service level, the mapping of applications or users to a particular service class, and the nature of the service delivered to each service class, need not be uniform across all routers nor stable over time.

From an individual session's point of view, the service still somewhat resembles best effort; sessions in a class are not insulated from each other, and there is no admission control to limit the number of sessions in the class.[5]

3. QUALITY DIFFERENTIATION FOR QOS

Quality of Service essentially means differentiating users or applications to service classes having different requirements from the network. This differentiation becomes a necessity when the network becomes overutilized. During times of congestion resources are divided in a manner which reflects most properly the QoS policy of an Internet Service Provider (ISP). The QoS policy may be built to offer similar quality for all or a set of users based on traffic classification rules. These traffic classification rules reflect the ISP's policy to justify some form of communication more 'socially valuable' than other.

Some of the methods for quality differentiation are presented in Figure 1. Basically the classification can be initiated either by the user or by the network. Within this division several possibilities and suggested methods exist.

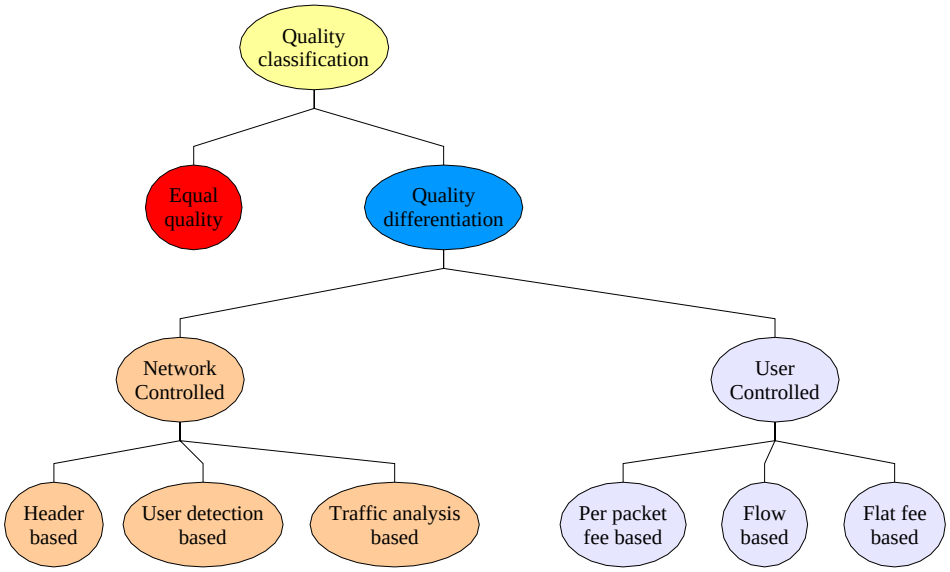

Figure 1. The ways of quality differentiation

3.1. Quality differentiation in DiffServ

In Differentiated Services the approach to quality differentiation is accomplished through code points (DSCP) which reflect the forwarding treatment a packet should have. This forwarding treatment is a set of behavioral rules manifesting relative importance of the class and congestion behavior inside a class. If these code points are set by the user, the mixture of traffic in different classes in the network is somewhat uncontrollable by the network provider. The ISP can remark packets entering the network, but the at that point the user loses control over the traffic mix. Network can also control the quality differentiation totally. Then user allows network to make decission on the best interest of user.

3.2. Quality differentiation in IntServ

Integrated Service model favors the per flow treatment of traffic, which is based on dynamic service level agreement (SLA). To set a dynamic SLAs signaling (RSVP) and state information inside the network is required. This concept creates a scaleability problem if there is a large number of users and applications in the network. This is an issue that has been seen as a major drawback for the Integrated Services approach.

3.3. Application perspective

From the functional perspective applications can be grouped to categories manifesting the reason of usage or nature of their operation. This grouping spreads over the transport protocols thus complicating management of traffic in individual classes.

Broad division of the applications can be done based on their functional behavior. This behavior is dependent on the purpose of the application and protocol used in communication. The key difference between classes is the issue of time. Time integrity is highly relevant for a large number of applications.

The functional grouping of application could be considered as follows:

- Interactive collaborative applications require constant attention from the user. These applications are used for human to human communication; like IPtelephony, video conversation, shared white board etc. They require minimum delay and delay variation through the network. The amount of bandwidth they use is relatively small from kilobits to tens of kilobits per second for voice applications and tens to hundreds of kilobits per second for video. These applications are usually implemented to use either raw UDP-protocol or UDP with RTP-like flow control for their communication.

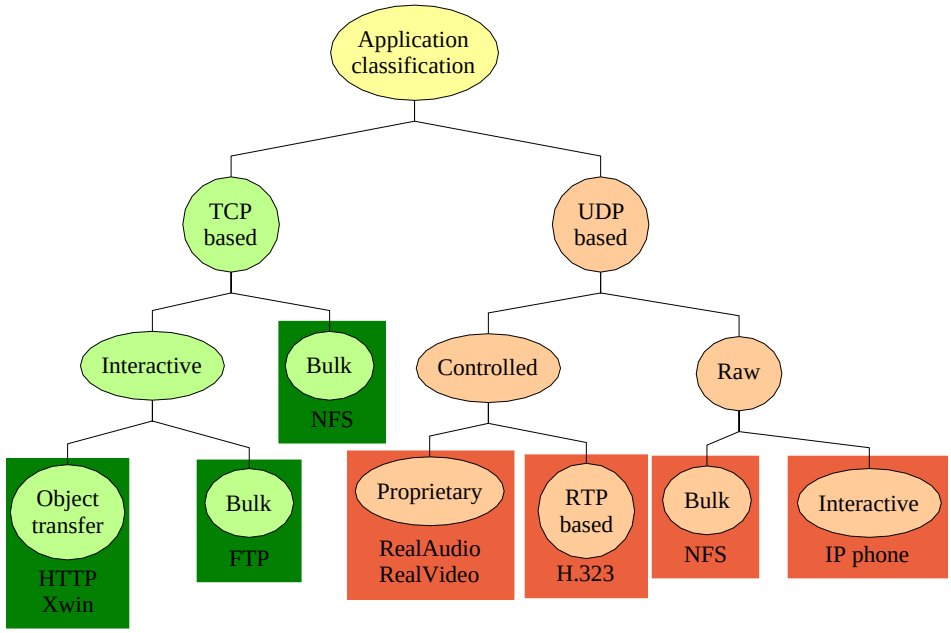

Figure 2. Division of link resources to classes of application

- Interactive applications form a large group of applications which can be categorized further to the streaming applications and object transfer applications.

  1. Streaming applications are continuous media playback applications like RealAudio and RealVideo. They produce traffic which benefits from bounded delay but the variation is allowed to disperse quite freely. Correctness of the information is much more relevant than the minimal transfer time. This is due to the jitter compensation they use. Jitter compensation algorithms usually estimate delay variation at the receiver and based on the result they either buffer more or less at the receiving end. These applications usually use the UDP-protocol enhanced with RTP or some proprietary mechanism to implement rate and error control.
  2. Object transfer applications perform tasks which are strongly related to the operating systems and moving of digital objects. This group consists of applications and protocols like X, NFS, HTTP and Telnet. The feedback which this group gives is directly observable from user so timely transfer is important. Rate control in this group is usually performed through the TCP flow and congestion control, so these applications can adjust their operation to the changing environment.

- Bulk transfer applications; FTP, email, nntp and other raw data transfer protocols belong to this class. These applications are used for transferring of large quantities information on the background. They use both TCP and UDP to perform tasks. They are not usually sensitive to time rather correctness of information is preferred.

4. USER CONTROLLED DIFFERENTIATION

User controlled differentiation means mechanisms by which the user can set quantitative or qualitative preferences to the traffic he is sending and/or receiving. These mechanisms are based on dynamic SLAs.

User controlled differentiation requires proper incentives to provide reasonable traffic mix. One idea is that the Internet should be operated as a free economy where supply and demand connect through the price. For a user there is an obvious objective to minimize monetary cost for a given quality of service, whereas for ISP pricing and the economics provide a tool to allocate scarce resources among competing users.

So in a sense price steers the differentiation in a user controlled case. It may be the only paramenter considered or it may set a reference point for some measureable parameter.

4.1. Flat fee differentiation

Flat fee is the prevalent charging scheme in the Internet. Flat fee is applied to cover costs of an user with an unlimited access to ISPs backbone. This unlimited access is truncated only by the access technology which commonly is PSTN/ISDN for the residential user and $nx64$ kbit/s leased line for the business user. This means that there is no quality related differentiation provided even though higher price is paid. The only differentiation point is the access capacity which may give better sustainable performance when more packets arrive to the queue from high speed access link compared to packets arriving from a slower speed access link duringthe times of congestion.

4.1.1. Flat fee for quality differentiation

Flat fee can also be used as a reference point for the quality differentiation. In Simple Integrated Media Access (SIMA)[6] the concept of flat fee is related to the resources the user is subscribing. Here, flat fee corresponds to Nominal Bit Rate (NBR), an expected user resource demand under normal quality operation, which acts as a reference point in the calculation for the actual quality level of the packet.

Quality in SIMA is implemented through priorities. Network has eight levels of priorities which each represent a quality level. Each network device decides individually whether a packet in a certain quality level should be accepted to the queue or discarded. The acceptance is based on the current workload of the device.

Users have the control over the quality by two different factors: in a long run he may purhace higher NBR or he may trade bandwidth for quality. By sending packets slower than the subscribed NBR, access device recognizes user who wishes to retrieve better quality. Based on the logarithmic ratio of the momentary bit rate (MBR) and NBR, a priority level for the packet is then decided. If lower qualities are acceptable user may send traffic at speeds higher than the NBR with the penalty that large amount of packets may very well be discarded during congestion.

4.2. Per packet fee differentiation

Per packet fee differentiation is based on the smart market mechanisms introduced in.[7–10] In smart market each packet carries a bid. In a congested router a bid is used in an auction to indicate how much the user is willing to pay for the successful delivery of the packet. What user actually pays for is the market-clearing price, which is the bid of the last accepted packet (lowest accepted bidding). This type of operation resembles space priority based operation and it does effectively allow higher priority traffic pass with lower loss probability and thus implements some quality differentiation. Bidding is necessary only during the times of congestion which makes bidding in some sense implementable without excessive ticketing effort.

4.3. Per flow differentiation

Finally, the per flow differentiation is the basic approach for the Integrated Services model. It requires incentives to guide the resource usage. The incentive here is that the per connection fee which is strongly related to connection oriented networking (PSTN/ISDN). In connection oriented networking resources are reserved for a single user for an unlimited time and with accumulating cost. This reservation lowers network ability to serve other users which is then experienced as flow and call blocking by other users.

5. NETWORK CONTROLLED DIFFERENTIATION

In the Differentiated Services -architecture it is also possible to offer pure network controlled differentiation. This means that the network is responsible for the separation of traffic to proper service classes. This separation is based on the ability of traffic conditioners in the access router to perform multi-field classification. Multi-field classification makes it possible to detect applications and users from the traffic stream. Broadly speaking, traffic classifiers are developed to assist network manager in separation of applications to appropriate service classes. Applications in general are defined based on the source port of the transmission protocol (TCP/UDP).

5.1. Header based differentiation

Header based differentiation is based on the analysis of the complete or partial packet header. Based on the extracted information some certainty of the application type and characteristics may exist, especially if application identifiers* are in the well-known port number area. Most commonly packet headers are examined to determine a class for a flow, the granularity of which is fivetuple: the source and destination IP-addresses, the source and destination ports, and the protocol identifier. This mechanism could also be used in conjunction with QoS-routing as means to direct application flow aggregates to capable links.

The combinations for filters to implement these classifiers are vast and therefore mechanisms to update and develop suitable filters are required. These mechanisms may be straight forward static lists of applications updated by network managers or dynamic application lists based on the results of applying some form of computational mechanisms to network measurement data. Some classifiers based on these kinds of mechanisms are introduced and evaluated in.[11–15]

5.2. User based differentiation

Quality differentiation can also be done based on the user or user group. Without any other requirements, an IP header may be processed in order to relate the source address and required treatment of the packet. To make this approach to work, the IP router must have a separate routing table entry for each user or user group, thus increasing the routing table size roughly by a factor of 10 to 100. This increase in routing table size causes more processing in router and thus slowers its ability to forward traffic. This problem can be partly solved by caching most frequent lookups and by implementing proper hash functions.[15–17]

Quality differentiation based on the identification of the user / user group may also be conducted as a side product of the access phase to an ISP network. Today ISPs implement several authentication mechanisms like RADIUS.[18] These authentication mechanisms try to prevent unauthorized use of ISPs resources. As a side product a policy database can be consulted and the resulting outcome can be downloaded to the router by using protocols like LDAP.[19] This type of dynamic user based policy networking is largely introduced by Directory Enabled Networking Ad Hoc Working Group in their proposal.[20]

There is a problem with a user detection based differentiation which does not lend itself to have an easy solution - the problem of different application demands. User detection based classification gives the same precedence for all traffic coming from the same user. This may be problematic if resources are scarce and at the same time there are high speed file transfers and interactive video applications competing from the bandwidth.

5.2.1. Per flow traffic analysis based differentiation

A per flow analysis of packets may be done in order to find out flows which could benefit from a higher level of quality. This per flow analysis may be based on the amount of information received on a certain channel,[11,21,13–15] the size of packets in a flow[22] or it may be based on the packet arrival pattern in a flow.[23]

- If the classification is based on the received amount of data and forwarding treatment is launched after a certain threshold then the system is favoring flows which send large amount of packets. The control of the actual quality differentiation outcome is minimal since there are relatively few parameters to tune for a desired outcome. In addition the misordering of packets may cause a problem if the threshold level is relatively high.
- Differentiation based on the packet lengths makes it possible to distinguish sessions having real-time requirements. Typical real-time applications use small packet sizes as explained in [24]. If the flow has continuously short packets it can be classified as interactive with great confidence.
- Classification mechanisms which use interarrival time histograms are theoretically interesting. By using interarrival times a firm division of applications into two classes can be made: (i) realtime and (ii) non-realtime. Realtime applications typically lack feedback control or at least they tend to keep constant intersending times of packets during the active period, talk spurt, for instance. This leads to a heavily uni-modal distribution with noise inserted as packet gaps during the silent periods. Non-realtime applications on the other hand rely heavily on window based flow control where a certain amount of data is sent to the network and whereafter an acknowledgment for the transmitted data is required. This kind of two phase operation will lead to heavily bi-modal distribution of interarrival times.[23]

*The application identifier here is the TCP/UDP port number

5.3. Traffic analysis based differentiation

Traffic analysis based differentiation means that background processors are used to build packet classifiers which most properly reflect the desired differentiation outcome. Differentiation is based on the statistical analysis of the traffic. Some classifiers of this kind are presented in.[12,25–27] These traffic classifiers are used to determine appropriate packet header filters for application based packet header classifier.

Traffic analysis is usually based on the aggregated flow analysis with the flow granularity of the fivetuple. Traffic is analyzed in dimensions, like the relative number of flows or the relative number of packets or a combination of the both in certain applications. Some form of classifier, like k-nearest neighbor, is needed to resolve to which class area individual data points belong to.

Due to the (usually) large amount of data points and relatively high uncertainty of behavior of some applications, artificial intelligence mechanism may be used. In [26] the use of Learning Vector Quantization (LVQ) is introduced to make classification decisions in a flow based connection oriented IP router environment. This approach is extended to have multiple application classes in [28]. Multiple application classes are resolved based on the observations of the nature of communication presented in Figure 2. To extract this type of information from packet streams different time-out values are used in the related flow analysis process. Short time-out values put the weight on the interactive flows whereas longer time-outs tend to weight also the less interactive communication. This approach could substantially benefit if further developed to take into account also the packet lengths. This would create a classifier which made possible to characterize interactive traffic based on packet length and inter-arrival times.

5.4. Problems of traffic differentiation

Several problems and unresolved issues exist in traffic differentiation:

1. One imminent problem with most of the differentiation models is the directionality. Directionality is the direction of packet flow where the action is taken. Without any state transfer from user access point to destination point(s) all except the application based differentiation concept are applied only to upstream traffic. This is recognized as a problem in the DiffServ community, since on the average only 10% traffic is flooded upstream and 90% floods downstream.

2. The second problem with the differentiation models, especially with the application based solutions, is reliability. This is a problem because of the possibility to change application identifiers (transport protocol port) to those that are widely used for some other form of communication. This may lead to a misclassifications.

3. The third problem is security. It is obvious that much of the communication may and is going to move for the use of some form of encryption. If mechanisms like SSH and IPSec are used the real nature of communication is hidden behind a single application identifier. This leads to the problem of classifying traffic to a real class when only parts of a packet are understandable. These problems are worst with application based differentiation.

Out of these problems we see that the first one is the hardest one. The whole idea of differentiating traffic and the issue of designing reasonable service models out of the traffic differentiation schemes is impossible if the control for classification is possible only to upstream. Problems two and three are crucial but they can be resolved by careful technical planning.

6. TRAFFIC CLASSIFICATION METHODS INVESTIGATED

Traffic classification is the process where a single stream of traffic is separated to logically different streams of traffic which by default have isolation between them inside a router. This separation is done by filtering the traffic stream. Filters are used to find pre-specified patterns from IP packets or on their arrival process. These patterns may reside in the header or the payload as long as the place is accurately defined. Typical locations for filter patterns are the source IP address, destination IP address, protocol ID, source port, destination port and/or packet length. Filter patterns may contain a single value, range of values or a masked value.

6.1. Aggregated traffic analysis approach

Aggregated traffic analysis is based on the background processing of measured information. Based on the analysis we generate as simple filters as possible using general filter patterns. Aggregated traffic analyzer can use the same mechanisms as the real-time classifiers with the extra possibility to perform time consuming correlation analysis. The goal in the design of an aggregated traffic classifier is to introduce a network service profile (NSP) which would maximize the user perceived quality - or that the network would show convergence to a situation where each of the QoS classes is loaded with traffic from proper and suitable applications.

In [26] classification is done based on the per flow analysis with packet/per flow ratios. Flows are constructed based on 'standard' flow definition of Claffy et al.

In [28] same model is extended to cover multiple levels classification by modifying the notion of flow. Modification is done by changing the time-out of a flow. Time-out is the maximum interval between two consecutive packets, with similar flow dependent identifiers, belonging to the same flow. Time-out was chosen to be vector with values of 100ms and 1s. This vector based analysis gives two parallel packets per flow results.

Both of these previous methods exhibit several drawbacks which are investigated here more closely. (i) There is a single class for UDP and TCP traffic and (ii) there is a set of applications which tend to show characteristics of interactive application in packet per flow analysis even though they don ot belong to this application class.

These problems can be overcome with slight modifications to previous methods. First modification would be logical: Traffic is analyzed and handled separately for TCP and UDP i.e. there should be parallel classes for TCP and UDP. Solution for the second problem is also trivial. Packets are filtered based on their length in flow analysis with proper low pass filter. These two modifications should be broad enough to cover all foreseeable modifications in traffic an application space.

6.2. Packet length approach

We present here results from the analysis of measurements done in FUNET and TCT backbone. These results, in Figure 3, show strong correlation, as expected, between application type and packet length distribution. This would suggest that proposed method of Cheng and Kung in [22,23] works if classes in general behave as expected based on these candidate applications.

We examine these conditions based on the raw statistical analysis of traffic. If most, more than 90%, of the packets in application port are shorter than threshold value, 128 and 256 bytes, application is classified as interactive. Based on these statistical analysis and observations from the Table 1 and 2 we claim that packet length is not adequote to catch general behavior of class. Large amount of TCP ACKs causes some irrelevant applications to be classified as interactive. Nevertheless packet length approach shows its usefulness as a first hand easily implementable classifier.

6.3. Interarrival time approach

This approach is based on the observation that many of the interactive applications have a uni-modal interarrival time distribution as opposed to the other types of distributions for the elastic applications. Classification based on the distribution requires some information to be passed on 'default class' before actual classification can be done. Comparison of measured distribution against class prototypes could be done by using pattern recognition.

Two alternative approcahes apply after this

- The traffic flow could be assigned to a service level after it has been recognized as a classified flow. This has the downside of letting some part of the flow enter the network without the classification decisions.
- The pattern recognition process would attach suitable application information (port numbers, for instance) after it has made the classification decision. After this the classification would, once again, return to using a selected set of port numbers as the NSP. This approach would avoid the transient service level assignments in the beginning of the flow.

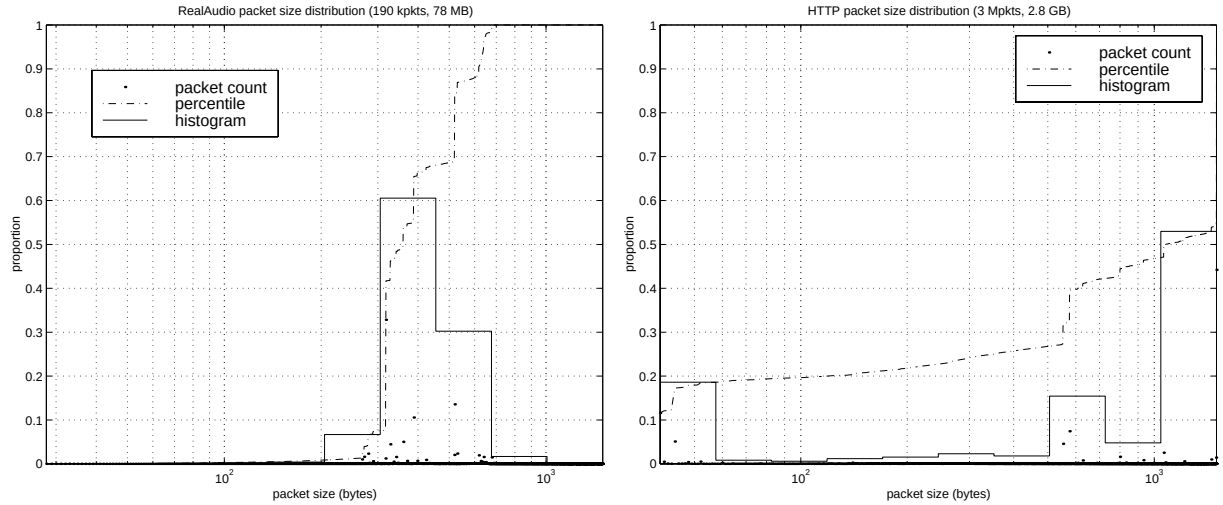

(a) RealAudio

(b) WWW

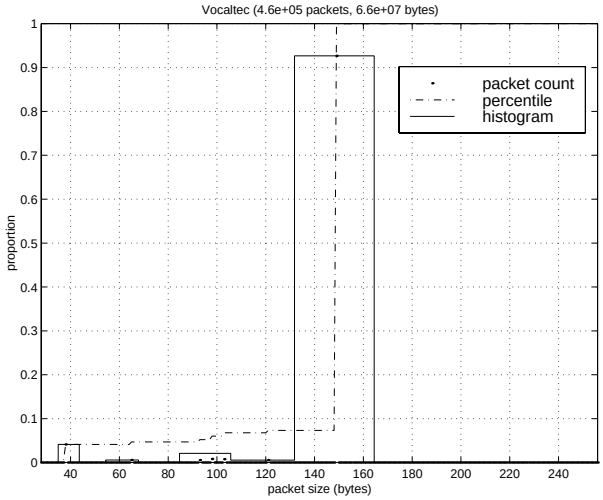

(c) Vocaltec

Figure 3. Packet length distribution of typical applications in different classes

7. CONCLUSIONS

We have presented several mechanisms to do quality differentiation and related classication of traffic to different classes. These mechanisms are all applicable in real world to do their tasks, some perform better in environment aimed to offer certain type of service than other. We have identified and more closely observed application based quality differentiation and classifiers which use realtime or near realtime measurements and calculations to do classification decisions. The reason for observing these in detail was the problem of traffic directionality which hinders most of the differentiation methods to operate in light overhead manner. We will conclude as a result of observations that application based quality differentiation is readily implementable and mechanisms to do classification are very well implementable with todays technology.

Table 1. Application identifiers which contain more than 90% packets shorter than 128 or 256 bytes. Trace is from FUNET backbone June 1999.

UDP		TCP			
ntp	123	ftp	21	simba−cs	1543
netbios−ns	137	smtp	25	mvel−lm	1574
blackjack	1025	auth	113	?	1848
iad2	1031	?	1019	?	2031
loaprobe	1634	?	1024	?	2191
combox−web−acc	2534	?	1038	?	2248
?	3206	?	1039	psbserver	2350
?	3223	?	1040	command−mq−gm	2664
icq	4000	?	1049	sns−channels	3380
dtspcd	6112	?	1050	?	5788
realaudio	6970	?	1051	?	6668
realaudio	6974	?	1055	?	6669
realaudio	7119	?	1088	?	7776
?	8153	?	1094	ndmp	10000
?	25651	rmiactivation	1098	?	27909
?	27001	?	1100		
?	27901	?	1113		
?	28800	?	1230		
?	62953	?	1253		
?	64592	?	1266		

(a) 128B

UDP		TCP			
domain	53	ftp	21	?	2031
ntp	123	telnet	23	?	2191
netbios−ns	137	smtp	25	?	2248
blackjack	1025	auth	113	psbserver	2350
iad2	1031	?	1019	command−mq−gm	2664
loaprobe	1634	?	1024	sns−channels	3380
combox−web−acc	2534	?	1038	?	5788
?	3206	?	1039	?	6667
?	3223	?	1040	?	6668
icq	4000	?	1049	?	6669
dtspcd	6112	?	1050	?	6767
realaudio	6970	?	1051	acmsoda	6969
realaudio	6974	?	1055	?	7776
realaudio	7119	?	1088	ndmp	10000
?	8153	?	1094	?	27909
snapenetio	22000	rmiactivation	1098		
?	25651	?	1100		
?	27001	?	1113		
?	27504	?	1230		
?	27901	?	1253		
?	27910	?	1266		
?	28800	simba−cs	1543		
?	62953	mvel−lm	1574		
?	64592	?	1848		

(b) 256B

Table 2. Application identifiers which contain more than 90% packets shorter than 128 or 256 bytes. Trace is from TCT backbone July 1999.

UDP		TCP			
ntp	123	smtp	25	?	1141
dns query	1025	?	1017	?	1142
?	1076	?	1019	?	1143
icq	4000	?	1020	?	1148
realaudio	6970	?	1021	?	1154
		?	1024	?	1173
		?	1082	?	1199
		ansoft−lm−2	1084	?	1203
		?	1089	?	1225
		sunclustermgr	1097	?	1229
		?	1105	?	1233
		?	1107	?	1251
		?	1115	?	1317
		?	1119	hiq	1410
		?	1121	innosys	1412
		?	1125	firefox	1689
		?	1140	?	5280

(a) 128B

UDP		TCP			
ntp	123	smtp	25	?	1141
dns query	1025	?	1017	?	1142
?	1076	?	1018	?	1143
icq	4000	?	1019	?	1148
realaudio	6970	?	1020	?	1154
vocaltec−phone	22555	?	1021	?	1173
		?	1024	?	1199
		?	1082	?	1203
		ansoft−lm−2	1084	?	1225
		?	1089	?	1229
		sunclustermgr	1097	?	1233
		?	1105	?	1251
		?	1107	?	1317
		?	1115	hiq	1410
		?	1119	innosys	1412
		?	1121	firefox	1689
		?	1125	?	5280
		?	1140	?	6000

(b) 256B

Acknowledgements

We would like to express our gratitude to the Finnish University and Research Network, FUNET, for providing valuable data traces of their backbone. Marko Luoma and Markus Peuhkuri were supported by Academy of Finland from the contract for project MI^{2}TTA. Mika Ilvesmäki is supported by the Nokia Foundation and Technology Research Center from the contract for project IPANA.

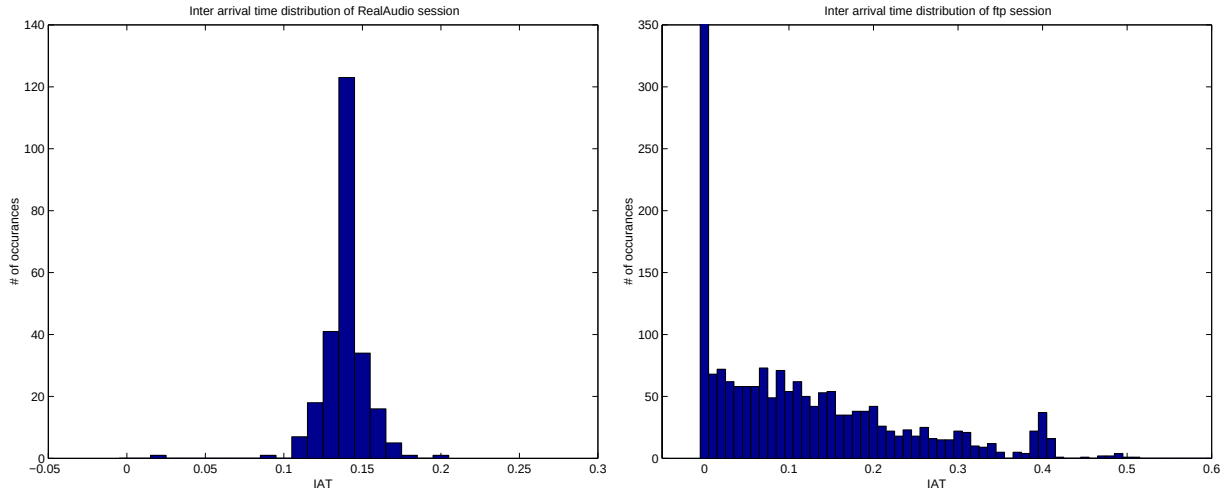

(a) RealAudio

(b) FTP

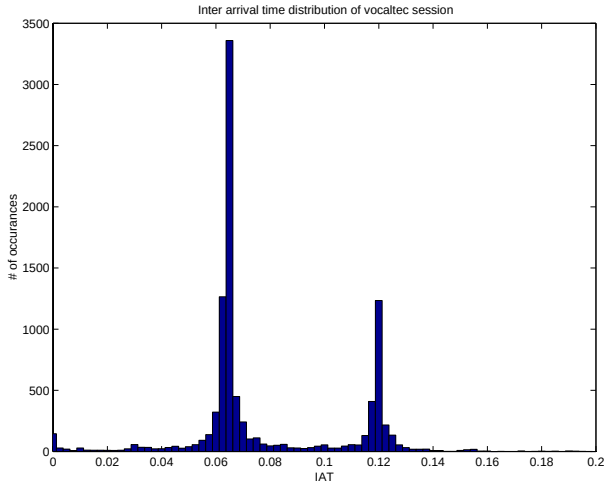

(c) Vocaltec IP-phone

Figure 4. Inter arrival time distribution of typical realtime application and elastic application.

REFERENCES

1. J. Wroclawski, "The use of rsvp with ietf integrated services," Tech. Rep. RFC 2210, IETF, September 1997.
2. J. Wroclawski, "Specification of the controlled-load network element service," Tech. Rep. RFC 2211, IETF, September 1997.
3. S. Shenker, C. Partridge, and R. Guerin, "Specification of guaranteed quality of service," Tech. Rep. RFC 2212, IETF, September 1997.
4. Y. Bernet, J. Binder, S. Blake, M. Carlso, S. Keshavn, E. Davies, B. Ohlman, D. Verma, Z. Wang, and W. Weiss, "A framework for differentiated services," Tech. Rep. draft-ietf-diffserv-framework-02.txt, IETF, February 1999.
5. M. Balakrishnan and R. Venkateswaran, "Qos and differentiated services in multiservice network environment," *Bell Labs Technical Journal* , pp. 222–238, October-December 1998.
6. K. Kilkki, "Simple integrated media access," Tech. Rep. draft-kalevi-simple-media-access-01.txt, IETF, June 1997.

7. J. MacKie-Mason and H. Varian, "Some economics of the internet," in *Networks, Infrastructure and the New Task for Regulation*, W. Sichel, ed., University of Michigan Press, 1996.
8. J. MacKie-Mason, "A smart market for resource reservation in a multiple quality of service information network." Preliminary Draft, September 1997.
9. J. MacKie-Mason and H. Varian, "Pricing congestible network resources," *IEEE Journal on Selected Areas in Communications* **13**, pp. 1141–1149, September 1995.
10. J. MacKie-Mason and H. Varian, "Pricing the internet," in *Public Access to the Internet*, B. Kahin and J. Keller, eds., pp. 269–314, MIT Press, 1995.
11. M. Ilvesmki, K. Kilkki, and M. Luoma, "Packets or ports - the decisions of ip switching.," SPIE Proceedings series, pp. 53–64, SPIE, 1997.
12. M. Ilvesmäki, R. Kantola, and M. Luoma, "Adaptive flow classification in ip switching: The measurement based approach," in *Internet Routing and Quality of Service*, R. Onvural, S. Civandlar, P. Doolan, and J. Luciani, eds., vol. 3529 of *SPIE Proceedings series*, pp. 277–286, SPIE, SPIE, November 1998.
13. P. Newman, G. Minshall, and T. L. Lyon, "Ip switching – atm under ip," *IEEE/ACM Transactions on Networking* **6**, pp. 117–129, April 1998.
14. S. Lin and N. McKeown, "A simulation study of IP switching," in *ACM SIGCOMM '97*, 1997.
15. P. Gupta and N. McKeown, "Packet classification on multiple fields," in *Proceedings of ACM SIGCOMM'99*, ACM, August 1999.
16. G. P. Chandranmenon and G. Varghese, "Trading packet headers for packet processing," *IEEE/ACM Transactions on Networking* **4**(2), pp. 141–151, 1996.
17. P. Gupta, S. Lin, and N. McKeown, "Routing lookups in hardware at memory access speeds," in *Proceedings of INFOCOM'98*, IEEE, April 1998.
18. C. Rigney, A. Rubens, W. Simpson, and S. Willens, "Remote authentication dial in user service (radius)," Tech. Rep. RFC 2138, IETF, April 1997.
19. M. Wahl, T. Howes, and S. Kille, "Lightweight directory access protocol (v3)," Tech. Rep. RFC 2251, IETF, December 1997.
20. S. Judd and J. Strassner, "Directory enabled networks - information model and base schema," Tech. Rep. Draft v3.0c5, Directory Enabled Networks Ad Hoc Working Group, 1998.
21. H. Che, S.-Q. Li, and A. Lin, "Adaptive resource management for flow-based ip/atm hybrid switching systems," *IEEE/ACM Transactions on Networking* **6**, pp. 544–557, October 1998.
22. B. Nandy, N. Seddigh, A. Chapman, and J. H. Salim, "A connectionless approach to providing qos in ip networks," in *High Performance Networking*, H. van As, ed., pp. 363–379, IFIP, Kluwer Academic Publishers, September 1998.
23. A. Chapman and H. Kung, "Automatic quality of service in ip networks," in *Proceedings of the Canadian Conference on Broadband Research*, pp. 184–189, April 1997.
24. K. van der Waal, M. Mandjes, and H. Bastiaansen, "Delay performance analysis of the new internet services with guaranteed qos," *IEEE Proceedings* **85**, pp. 1947–1957, December 1997.
25. J. Karvo and M. Ilvesmäki, "Nondeterministic classifier performance evaluation for flow based ip switching," in *High Performance Networking*, H. van As, ed., pp. 613–624, IFIP, Kluwer Academic Publishers, September 1998.
26. M. Ilvesmäki, M. Luoma, and R. Kantola, "Flow classification schemes in traffic-based multilayer ip switching - comparison between conventional and neural approach," *Computer Communications* **21**(13), pp. 1184–1194, 1998.
27. M. Ilvesmäki, M. Luoma, and R. Kantola, "Learning vector quantization in flow classification of ip switched networks," in *Proceedings of GLOBECOM'98*, vol. 5, pp. 3017–3022, IEEE, November 1998.
28. M. Ilvesmäki and M. Luoma, "Performance analysis of multi-class internet traffic classifier in a connection oriented router environment," in *Submitted to SPIE Voice, Video and Data Communications 1999*, 1999.