

A Method for Evaluation of Quality of Service in Computer Networks

Tomasz Bujlow, Sara Ligaard Nørgaard Hald, Tahir Riaz, Jens Myrup Pedersen

Section for Networking and Security, Department of Electronic Systems

Aalborg University, DK-9220, Aalborg East, Denmark

tbu@es.aau.dk, slh@es.aau.dk, tahir@es.aau.dk, jens@es.aau.dk

Abstract—Monitoring of the Quality of Service (QoS) in high-speed Internet infrastructures is a challenging task. However, precise assessments must take into account the fact that the requirements for the given quality level are service-dependent. The backbone QoS monitoring and analysis requires processing of large amounts of data and the knowledge about the kinds of applications, which generate the traffic. To overcome the drawbacks of existing methods for traffic classification, we proposed and evaluated a centralized solution based on the C5.0 Machine Learning Algorithm (MLA) and decision rules. The first task was to collect and to provide to C5.0 high-quality training data divided into groups, which correspond to different types of applications. It was found that the currently existing means of collecting data (classification by ports, Deep Packet Inspection, statistical classification, public data sources) are not sufficient and they do not comply with the required standards. We developed a new system to collect the training data, in which the major role is performed by volunteers. Client applications installed on volunteers' computers collect the detailed data about each flow passing through the network interface, together with the application name taken from the description of system sockets. This paper proposes a new method for measuring the level of Quality of Service in broadband networks. It is based on our Volunteer-Based System to collect the training data, Machine Learning Algorithms to generate the classification rules and the application-specific rules for assessing the QoS level. We combine both passive and active monitoring technologies. The paper evaluates different possibilities of the implementation, presents the current implementation of the particular parts of the system, their initial runs and the obtained results, highlighting parts relevant from the QoS point of view.

Index Terms—broadband networks, data collecting, Machine Learning Algorithms, performance monitoring, Quality of Service, traffic classification, volunteer-based system.

Manuscript received August 1, 2012. This research work was conducted in the context of a PhD project in the Section for Networking and Security in the Department of Electronic Systems at Aalborg University. The project is co-financed by the European Regional Development Fund (ERDF) – see http://ec.europa.eu/regional_policy/thefunds/regional/index_en.cfm, and Bredbånd Nord A/S, a regional fiber networks deliverer (see <http://www.bredbaandnord.dk/>).

Tomasz Bujlow is with Section for Networking and Security, Department of Electronic Systems, Aalborg University, Aalborg East, 9220 Denmark (tel: +45 5026 2494, e-mail: tbu@es.aau.dk).

Sara Ligaard Nørgaard Hald is with Section for Networking and Security, Department of Electronic Systems, Aalborg University, Aalborg East, 9220 Denmark (e-mail: slh@es.aau.dk).

Tahir Riaz is with Section for Networking and Security, Department of Electronic Systems, Aalborg University, Aalborg East, 9220 Denmark (e-mail: tahir@es.aau.dk).

Jens Myrup Pedersen is with Section for Networking and Security, Department of Electronic Systems, Aalborg University, Aalborg East, 9220 Denmark (e-mail: jens@es.aau.dk).

I. INTRODUCTION

This journal paper is an extended and revised version of [1], which was presented at the 14th International Conference on Advanced Communication Technology (ICACT 2012).

One of the most interesting challenges in today's world is how to measure the performance of computer network infrastructures, when different types of networks are merged together. In the last few years, the data-oriented networks evolved into converged structures, in which the real-time traffic, like voice calls or video conferences, is more and more important. The structure is composed of traditional data cable or more modern fiber links, existing Plain Old Telephone Service (POTS) lines used to provide analog services (voice telephony), or digital services (ADSL, PBX, ISDN), and nowadays also of mobile and wireless networks. There are numerous methods for the measurement of Quality of Service (QoS) in the current use, which provide the measurements both on the user side and in the core of the network. Internet Service Providers are interested in centralized measurements and detecting problems with particular customers before the customers start complaining about the problems, and if possible, before the problems are even noticed by the customers.

Each network carries data for numerous different kinds of applications. QoS requirements are dependent on the service. The main service-specific parameters are bandwidth, delay, jitter, and packet loss. Regarding delay, we can distinguish strict real time constraints for voice and video conferences, and interactive services from delivery in relaxed time frame. In a conversation, delay of about 100 ms is hardly noticeable, but 250 ms of delay means an essential degradation of the transmission quality, and more than 400 ms is considered as severely disturbing [2].

Therefore, in order to provide detailed information about the quality level for the given service in the core of the network, we need to know, what kinds of data are flowing in the network at the present time. Processing all the packets flowing in a high-speed network and examining their payload to get the application name is a very hard task, involving large amounts of processing power and storage capacity. Furthermore, numerous privacy and confidentiality issues can arise. A solution for this problem can be the use of Machine Learning Algorithms (MLAs), which use previously generated

decision rules, which are based on some statistical information about the traffic. In our research, we used one of the newest MLAs – C5.0. MLAs need very precise training sets to learn how to accurately classify the data, so the first issue to be solved was to find a way to collect high-quality training statistics.

In order to collect the necessary statistics and generate the training sets for C5.0, a new system was developed, in which the major role is performed by volunteers. Client applications installed on their computers collect the detailed information about each flow passing through the network interface, together with the application name taken from the description of the system sockets. The information about each packet belonging to the flow is also collected. Our volunteer-based system guarantees precise and detailed data sets about the network traffic. These data sets can be successfully used to generate statistics used as the input to train MLAs and to generate accurate decision rules.

The knowledge about the kind of application to which the traffic belongs obtained from MLAs can be used together with traffic requirements for the given application to assess the QoS level in the core of the real network. The real traffic needs to be sampled to obtain the necessary raw statistics. Parameters like jitter, burstiness, download and upload speed (and delay and packet loss for TCP traffic) can be assessed directly on the basis of the information obtained from the captured traffic. To assess the delay and packet loss for UDP traffic, active measurement techniques must be involved (like ping measurements in both directions).

The remainder of this document is split into several sections, which describe in detail the system architecture and some parts of the implementation. Section II contains the overview of the current methods for assessing the network QoS level. Both passive and active methods are described along with their advantages and weaknesses. Section III gives the overview of our methods, so the reader is able to understand how the particular components are built and connected with each other. Section IV describes the current methods used for traffic classification in computer networks and it explains why they are not sufficient for our needs. Section V presents our new tool used for collecting and classification of the network traffic – the Volunteer-Based System (VBS). Section VI shows how the statistical parameters are obtained from the data collected by VBS. Section VII evaluates different Machine Learning Algorithms and shows why we chose C5.0 to be included in our system. Section VIII demonstrates the design and implementation of the system, while Section IX summarizes the most important points.

II. RELATED WORK

During the last 20 years we have been witnesses to the subsequent and increasing growth of the global Internet and the network technology in general. The broadband and mobile broadband performance today is mainly measured and monitored by speed. However, there are several other parameters, which are important for critical business and real-time applications, such as voice and video applications or first-

person shooter games. These parameters include round trip time, jitter, packet loss, and availability [3], [4].

The lack of the centralized administration makes it difficult to impose a common measurement infrastructure or protocol. For example, the deployment of active testing devices throughout the Internet would require a separate arrangement with each service provider [3]. This state of affairs led to some attempts to make simulation systems representing real characteristics of the traffic in the network. Routers and traffic analyzers provide passive single-point measurements. They do not measure the performance directly, but the traffic characteristics are strongly correlated with the performance. Routers and switches usually feature a capability to mirror incoming traffic to a specific port, where a traffic meter can be attached. The main difficulty in passive traffic monitoring is the steadily increasing rate of transmission links (10 or 100 GB/s), which can simply overwhelm routers or traffic analyzers, which try to process packets. It forces the introduction of packet sampling techniques and, therefore, it also introduces the possibility of inaccuracies. Even at 1 Gbit/s, the measurements can result in enormous amounts of data to process and store within the monitoring period [3].

To overcome the heavy load in the backbone and to not introduce inaccuracies, a smart monitoring algorithm was needed. There are several approaches to estimate which traffic flows need to be sampled. A path anomaly detection algorithm was proposed in [5]. The objective was to identify the paths, whose delay exceeds their threshold, without calculating delays for all paths. Path anomalies are typically rare events, and for the most part, the system operates normally, so there is no need to continuously compute delays for all the paths, wasting processor, memory, and storage resources [5]. Authors propose a sampling-based heuristic to compute a small set of paths to monitor, reducing monitoring overhead by nearly 50 % comparing to monitoring all the existing paths.

The next proposals on how to sample network traffic in an efficient way were made on the basis of adaptive statistical sampling techniques, and they are presented in [6] and [7].

If a congestion is detected, from user's perspective it is very important to know, if the congestion is located on the local or on the remote side. If the link experiences a local congestion, the user may be able to perform certain actions, e.g. shut down an application, which consumes a lot of bandwidth, to ease the congestion. On the other hand, if the congested link is a remote link, either in the Internet core or at the server side, the back-off of the low-priority applications on the user's side is unnecessary. It only benefits the high-priority flows from other users, which compete for that link. Since this altruistic behavior is not desirable, the low priority TCP only needs to back off, when the congested link is local [8].

Detecting the location of the congestion is a challenging problem due to several reasons. First of all, we cannot send many probing packets, because it causes too much overhead, and it even expands the congestion. Secondly, without a router support, the only related signals to the end applications are packet losses and delays. If the packet losses were completely synchronized (packets were dropped from all the flows), the problem would be trivial. In the reality, the packet loss pattern

is only partially synchronized [8]. Authors of [8] attempted to solve the problem of detecting the location of the congestion by using the synchronization of the behavior of loss and delay across multiple TCP sessions in the area controlled by the same local gateway. If many flows see a synchronized congestion, the local link is the congested link. If the congested link is remote, it is less likely that many flows from the same host pass the same congested link at the same time. If there is only a small number of flows which see the congestion, the authors performed an algorithm based on queuing delay patterns. If the local link is congested, most flows typically experience high delays at a similar level. Otherwise, the congestion is remote [8].

The traffic can be profiled according to the protocol composition. Usually, the predominance of the TCP traffic is observed (around 95 % of the traffic mix). When a congestion occurs, TCP sources respond by reducing their offered load, whereas UDP sources do not. It results in the higher ratio of UDP to TCP traffic. If the proportion becomes high and the bandwidth available to TCP connections becomes too low to maintain a reasonable transmission window, the packet loss increases dramatically (and TCP flows become dominated by retransmission timeouts) [3]. Packet sizes provide insight into the types of packets, e.g. short 40-44 bytes packets are usually TCP acknowledgments or TCP control segments (SYN, FIN or RST) [3].

Active methods for QoS monitoring raise three major concerns. First, the introduction of the test traffic will increase the network load, which can be viewed as an overhead cost for active methods. Second, the test traffic can affect measurements. Third, the traffic entering an ISP can be considered as invasive and discarded or assigned to a low-priority class [3].

Within an administrative domain (but not across the entire Internet), the performance can be actively monitored using the data-link layer protocol below IP, as the Operations, Administration and Maintenance (OAM) procedure in ATM and MPLS networks. As a result, at the IP layer it is often desirable to measure performance using the IP/ICMP protocol. So far, most tools or methods are based on ping (ICMP echo request and echo reply messages) or traceroute (which exploits the TTL field in the header of the IP packet) [3].

Although the round-trip times measured by ping are important, ping is unable to measure the one-way delay without additional means like GPS to synchronize clocks at the source and destination hosts. Another difficulty is that pings are often discarded or low-prioritized by many ISP in their networks. Traceroute will not encounter this problem because UDP packets are used. However, traceroute has known limitations. For example, successive UDP packets sent by traceroute are not guaranteed to follow the same path. Also, the returned ICMP message may not follow the same path as the UDP packet that triggered it [3].

Although the end-to-end performance measurements can be carried out at the IP layer or the transport/application layer, the latest is capable of measurements closer to user's perspective. The basic idea is to run a program emulating a particular application that will send traffic through the Internet. All the

parameters (delay, packet loss, throughput, etc) are measured on the test traffic. This approach has one major drawback - a custom software needs to be installed on the measurement hosts [3].

On the basis of the mentioned work, we found out that the existing solutions are not sufficient for precise QoS measurements. This state of affairs motivated us to create a new system which combines both passive and active measurement technologies.

III. OVERVIEW OF THE METHODS

The flow chart of our system is shown in Figure 1. The following paragraphs contain the detailed description of our methods. At first, the volunteers must be recruited from the network users. The volunteers install on their computer a client program, which captures the relevant information about the traffic and submits the data to the server. On the server, these data are used to generate per-application traffic statistics. The C5.0 Machine Learning Algorithm uses these statistics to learn how to distinguish between different types of applications and, later, it generates the classification rules (decision trees).

In order to assess the network QoS level in the core of the network for particular users, we needed to find a method to capture the relevant traffic. The challenging task is to process significant amounts of traffic in high-speed networks. When the relevant flows are captured, per-flow statistics need to be generated. There are two kinds of statistics generated at this step: one used for determining the kind of application associated with that flow, and one used for assessing the QoS level in the passive way. The system uses previously generated classification rules together with the first type of statistics to find out which application the flow belongs to. Then, on the basis of the kind of the application, the system determines the ranges of values of the relevant QoS parameters. The last step is to check if the current values (obtained from the flow statistics or in the active way) match the expected ones. If not, the quality of the given service is considered as degraded.

IV. CURRENT METHODS FOR OBTAINING PRE-CLASSIFIED DATA

There are many existing methods for obtaining pre-classified data, but none of them were feasible to deliver the data required by us to obtain accurate statistics, which could be used to train Machine Learning Algorithms (MLAs). The traffic classification requires the packets to be logically grouped into some structures, which could be assigned to the particular application. The most common used structure among the classification methods is the flow defined as a group of packets, which have the same end IP addresses, ports, and use the same transport layer protocol. In this section, we describe the methods and evaluate their usefulness in providing the data to our system.

A. Capturing Raw Data from the Network Interfaces

The first possibility is to install one application at a time on a host, and to capture its traffic by an external tool, such as Wireshark [9]. Unfortunately, this approach is very slow

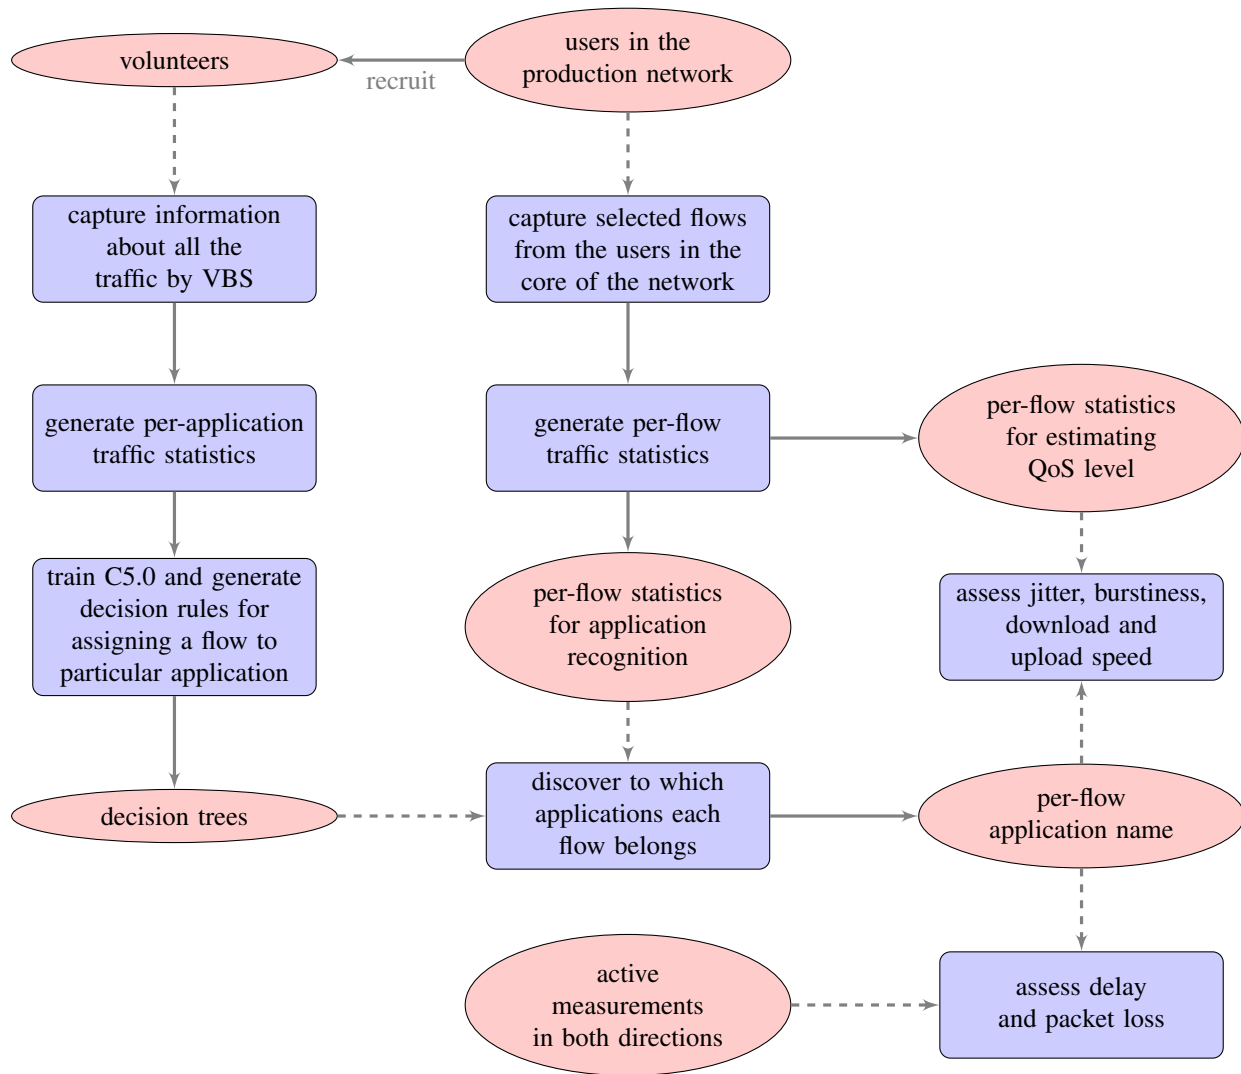


Figure 1. The flow chart of the system

and it is not scalable. At first, it requires us to install on a host each application that generates the traffic we want to capture. Before installing the application, we must uninstall all other applications that can generate any network traffic. The next drawback is that every operating system has some background processes and many of them transmit some data through the network. An example of such a process is the system updater, which can run in background. There is no simple way to recognize packets belonging to the traffic generated by the application intentionally run by us, so the captured sample contains variable percentage of noise. Finally, some applications, for example, web browsers, can generate various types of traffic. Raw traffic capturers cannot distinguish interactive web traffic, web radio podcasts, video transmissions, or downloads of big files, performed by the same browser.

B. Classification by Ports

The port-based classification [11], [12] is very fast, and it is supported on almost all the layer-3 devices in computer

networks. Unfortunately, this method is limited to services, protocols, and applications, which use fixed port numbers. It means that with big probability we can correctly classify, for example, traffic generated by e-mail clients and file transfer clients using File Transfer Protocol (FTP), when they use the default ports to connect to servers. However, even in this case we have false positives and false negatives. False negatives result from non-standard ports used in this example by SMTP, POP3, or FTP servers. When a network administrator changes the port used by the given service (due to security reasons), the traffic is not classified correctly. False positives result from malicious applications, which intentionally use some well-known port numbers to be treated in the network with a priority, or to be able to transmit data at all. Such situation exists when a Torrent user runs his client on port 80, which causes the traffic to be treated as if it originated from a web server. Another big concern of port-based classification is the inability of recognizing different types of traffic using the same transport-layer protocol and the same transport-layer port. This drawback is strongly visible in the example of HTTP




Figure 2. Overview of the VBS system [10]

traffic, which can consist of data generated by interactive web browsing, audio and video streaming, file downloads, and HTTP tunneling for other protocols. Finally, the classification made by ports is unable to deal with protocols using dynamic port numbers, like BitTorrent or Skype [9], [13], [14].

C. Deep Packet Inspection (DPI)

The big advantage of the Deep Packet Inspection (DPI) [15] is the possibility to inspect the content of the traffic. It includes both inspecting particular packets, and inspecting flows in the network as the whole. For that reason, that makes possible to distinguish different kinds of content generated by the same application, or using the same application-layer protocol, such as HTTP. However, DPI is slow and requires a lot of processing power [9], [13]. Therefore, due to high load in today's network infrastructures, it is not feasible to run DPI in the core of the network. The speed of the Internet connections provided to private users tends to increase much faster than the processing power of their machines, so performing DPI on user's machines became impossible in our case. Feasibility to perform DPI on the user side does not depend only on possessing the necessary processing power, but also on the user's impression. High CPU usage tends to slow down the machine and it causes additional side-effects, for example, a howling CPU fan. For that reason, full DPI can be done only in a limited number of cases, namely on fast machines using a slow Internet connection. DPI also brings privacy and confidentiality issues, as it can reveal some highly sensitive personal data, such as information about used credit cards, logins and passwords, websites visited by the user, etc [9]. Moreover, DPI is unable to inspect encrypted traffic. Finally, DPI depends on signatures of various protocols, services, and applications, which need to be kept up to date.

D. Statistical Classification

Solutions using statistical classification became quite popular during the last few years [15]. To its characteristics we can include fast processing and low resource usage. Statistical classifiers are usually based on rules, which are automatically generated from samples of data. Therefore, such kinds of

classifiers often make use of Machine Learning Algorithms (MLAs). Apart from all these advantages, statistical classifiers have one big drawback – they need to be trained on the samples of data. So the technique assumes that we have already correctly classified data, which we can provide as the input to train the statistical classifier. For that reason, we cannot use this method to collect and classify the initial portion of data.

V. VOLUNTEER-BASED SYSTEM

The drawbacks of the existing methods for the classification of traffic in computer networks led us to the conclusion that we need to design and build another solution. Therefore, we decided to develop a system based on volunteers, which captures the traffic from their network interfaces, and groups the traffic into flows associated with the application name taken from Windows or Linux sockets. The architecture and the prototype were described and analyzed in [16] and [10], and the first version of our current implementation was presented in [17]. Afterwards, the system was extended to support recognizing different kinds of HTTP traffic, and it was named Volunteer-Based System (VBS). The detailed description and evaluation of the extended version of VBS can be found in [18]. We released the system under *The GNU General Public License v3.0*, and we published it as a SourceForge project. The project website [19] contains all the information needed to use the system (binary packages, screenshots, documentation and bug tracking system) as well as to perform further development (source code, roadmap, comprehensive documentation of the source code).

The architecture of the system is shown in Figure 2. This cross-platform solution consists of clients installed on users' computers (Microsoft Windows XP and newer and Linux are supported), and of a server responsible for storing the collected data. The client registers information about each flow passing the Network Interface Card (NIC), with the exception of the traffic to and from the local network. The captured information are: start time of the flow, anonymized identifiers of the local and the remote IP addresses, local and remote ports, transport layer protocol, anonymized identifier of the global IP address of the client, name of the application, and

identifier of the client associated with the flow. The system also collects information about all the packets associated with each flow: identifier of the flow to which the packet belongs, direction, size, TCP flags, relative timestamp to the previous packet in the flow, and information about the HTTP content carried by the packet. It is worth mentioning that one flow can contain many higher-layer streams, for example, one TCP flow can contain multiple HTTP conversations. Each of these conversations can transfer different kinds of content, like web pages, audio and video streams, or downloads of big files. For that reason, we extract from HTTP headers the information necessary to precisely separate the HTTP streams, and we append the information about the type of the stream to the first packet of the stream.

The collected information is then transmitted to the server, which stores all the data in a MySQL database for further analysis. The system was shown in [18] to be feasible and capable of providing detailed per-application information about the network traffic. An example of the stored flows on the server side is shown in Table I. The IP addresses for privacy reasons are translated by a one-way hash function and they are stored as anonymized identifiers. The information about the packets belonging to one complete TCP conversation is presented in Table II. As shown, this is an HTTP communication, during which there were transferred two files of the same type with identifier 22 (*text/html*).

The data collected during our experiments by the Volunteer-Based System were used for training the C5.0 Machine Learning Algorithm to be able to recognize the traffic generated by different types of applications and different types of traffic. The first approach, focusing on distinguishing 7 different applications (Skype, FTP, torrent, web browser, web radio, America’s Army and SSH) and achieving accuracy of over 99% was described and evaluated in [20]. The second approach, focusing on recognizing different kinds of HTTP content (audio, video, file downloads, interactive websites) was presented in [21].

VI. OBTAINING PER-APPLICATION STATISTICS

The next step was to obtain the statistical profiles of flows for different applications. Therefore, we developed a tool for calculating statistics of several traffic attributes for each flow in the database, which fulfills our requirements. The statistics include 32 attributes based on sizes and 10 protocol-dependent attributes [20]. We suspect that the attributes based on sizes are independent of the current conditions in the network (like for example congestion). All the protocol-dependent attributes are very general. Precise port numbers are not used, but only the information about whether the port is well-known or dynamic. This way we avoid constructing a port-based classifier, but we can retain the information if the application model is more like client-server or peer-to-peer.

The general calculated statistics are [20]:

- Number of inbound / outbound / total payload bytes in the sample
- Proportion of inbound to outbound data packets / payload bytes

- Mean, minimum, maximum first quartile, median, third quartile, standard deviation of inbound / outbound / total payload size in the probe
- Ratio of small inbound data packets containing 50 B payload or less to all inbound data packets
- Ratio of small outbound data packets containing 50 B payload or less to all outbound data packets
- Ratio of all small data packets containing 50 B payload or less to all data packets
- Ratio of large inbound data packets containing 1300 B payload or more to all inbound data packets
- Ratio of large outbound data packets containing 1300 B payload or more to all outbound data packets
- Ratio of all large data packets containing 1300 B payload or more to all data packets
- Application: skype, ftp, torrent, web, web_radio, game, ssh

The protocol-dependent attributes are [20]:

- Transport protocol: TCP, UDP
- Local port: well-known, dynamic
- Remote port: well-known, dynamic
- Number of ACK / PSH flags for the inbound / outbound direction: continuous
- Proportion of inbound packets without payload to inbound packets: continuous
- Proportion of outbound packets without payload to outbound packets: continuous
- Proportion of packets without payload to all the packets: continuous

The precise process of obtaining these statistics was described in detail and evaluated in [20]:

VII. MACHINE LEARNING ALGORITHMS

In the recent literature, we can find numerous approaches to use Machine Learning Algorithms to classify the traffic in computer networks. The most widely used MLA classifiers are C4.5 [9] and its modified Java implementation called J48 [13], [22]. Based on statistical analysis, MLAs have the ability to assign a particular class (like P2P) even to traffic generated by unknown applications [9]. It was also proved in [22] that the statistical parameters for the encrypted and unencrypted traffic produced by the same application are similar and, therefore, the encrypted payload does not influence results of the training or the classification. The accuracy of the classification by MLAs was claimed to be over 95% [9], [11], [12], [14], [15], [23]–[25]. The analysis of the related work can be found in [20].

It was found in [12] that the results of the classification are most accurate when the classifier was trained in the same network as the classification process was performed. This may be due to different parameters, which are constant in the particular network, but which differ among various networks. A good example is the Maximum Transmission Unit, which can easily influence statistics based on sizes. Therefore, in our design, we decided to train the classifier by volunteers in the same network as the classifier will be installed. This allows us to make a self-learning system, where a group of volunteers

Table I
EXAMPLE OF THE STORED FLOWS DATA

flow id	client id	start time	local IP	remote IP	local port	remote port	protocol name	global client IP	application name
1193430	4	1325445237826039	d1e0229	fb70266	48293	25395	UDP	178a02f1	uTorrent
2393417	5	1325445237826176	f4c025e	12230296	2276	80	TCP	177d02ef	chrome
1193423	1	1325445237826304	d20022b	11920285	53778	80	TCP	12350297	firefox
1484673	4	1325445237825884	d1e0229	12170293	58104	993	TCP	178a02f1	thebat
3429674	4	1325445236820017	d1e0229	14cb02b9	61159	80	TCP	178a02f1	Dropbox
3329860	1	1325445237044777	d20022b	1199028a	47801	80	TCP	12350297	plugin-container
3829589	1	1325445236797638	d20022b	124d0296	36868	80	TCP	12350297	wget
3474027	4	1325445212663601	d1e0229	14db02c2	63409	24536	UDP	178a02f1	Skype
4194793	1	1325445280781252	d20022b	1206028f	53331	22849	TCP	12350297	amule

Table II
ONE TCP COMMUNICATION STORED IN THE DATABASE

flow id	direction	packet size [B]	SYN flag	ACK flag	PSH flag	FIN flag	RST flag	CWR flag	ECN flag	URG flag	relative timestamp [μ s]	content type
2784673	OUT	60	1	0	0	0	0	0	0	0	0	1
2784673	IN	60	1	1	0	0	0	0	0	0	30012	1
2784673	OUT	52	0	1	0	0	0	0	0	0	44	1
2784673	OUT	431	0	1	1	0	0	0	0	0	395	1
2784673	IN	52	0	1	0	0	0	0	0	0	30241	1
2784673	IN	527	0	1	1	0	0	0	0	0	2554	22
2784673	OUT	52	0	1	0	0	0	0	0	0	27	1
2784673	IN	539	0	1	1	0	0	0	0	0	10455	22
2784673	OUT	52	0	1	0	0	0	0	0	0	15	1
2784673	OUT	287	0	1	1	0	0	0	0	0	1383	1
2784673	OUT	52	0	1	0	1	0	0	0	0	15047	1
2784673	IN	269	0	1	1	0	0	0	0	0	16408	1
2784673	OUT	40	0	0	0	0	1	0	0	0	45	1
2784673	IN	52	0	1	0	1	0	0	0	0	13354	1
2784673	OUT	40	0	0	0	0	1	0	0	0	29	1

in the network delivers the data used for training the classifier constantly improving its accuracy, while all the users can be monitored in the core using the generated decision rules. The next advantage of the design is that even if some network users cannot participate in the data collecting process because of using other operating systems or devices than supported (like MacOS, Apple or Android smartphones), they will still be able to be monitored in the core of the network because of rules created on the basis of the data collected from the other users.

Our system uses the C5.0 MLA, which is a successor of C4.5. It is proven to have many advantages over its predecessor, such as higher accuracy, possibilities to use boosting, pruning, weighting and winnowing attributes. Furthermore, the time needed to generate the decision tree or rules drastically decreased [26]. In order to test the efficiency of C5.0, we performed a set of tests during which we used various training and classification options. The training statistics were obtained from the data provided by our VBS. During our research, we found relevant the set of arguments and discovered that the best results were obtained using the boosted classifier. The average accuracy fluctuated between 99.3% and 99.9%, depending on the number of training and test cases and the amount of data from each case. This behavior is illustrated in Figure 3. It is worth mentioning that in our experiment we considered only 7 different groups of applications and only flows longer than 15 packets. In our




Figure 3. Average error rates of the classifiers [20]

small-scale prototype for tests, we decided to limit the number of applications and take into account Skype, FTP, torrent, web traffic, web radio traffic, interactive game traffic, and SSH [20]. The limitation of the flow length was done because we needed to have at least 5 packets to generate the statistics (the first 10 packets of each flow were skipped as their behavior is different than the behavior of the rest of the flow). The detailed description of our methods and results can be found in [20]. The decision tree generated in this step can be used to classify the traffic in the real network.

VIII. CENTRALIZED MONITORING SOLUTION

This paragraph presents the proposed design of the centralized monitoring solution, which can be placed in any point in the network to examine the network QoS.

Because of the heavy load in the high-speed networks, it is not possible to monitor all the flows passing the central point at the same time. Therefore, only the statistics from selected flows can be captured and passed to C5.0. The selection of such flows can be based on two methods: capturing one flow per user and intelligent switching between the flows. From the QoS point of view, it is important to discover the problems with a particular user or to inform the user that the problems experienced by him are the result of problems in the remote network. If it is the user who has the problem, then the problem usually influences all the user's network activity.

Each application has some special requirements regarding the network parameters. When a small congestion occurs, the service level can still be sufficient for P2P file downloads, but Skype communication may be not possible because of big jitter and delays. It is, therefore, not sufficient to monitor one random flow at a time, but we need to monitor a flow which have high quality requirements. Our solution should be built based on the following assumptions:

- Only one flow per user at a time is consistently monitored for QoS.
- Statistics for another random flow per user at a time are passed to C5.0 to discover the application.
- If the application has higher QoS requirements than the currently monitored, switch monitoring to the new flow; if not, stick to the current one.
- If the monitoring of the selected flow discovers problems, start monitoring a few flows at a time to check if this problem lay on the user's side or on the remote side.

Because of the dynamic switching between the flows when determining the application, it is most probable that the system will not be able to capture flows from their beginning. The classifier designed by us, which uses C5.0, is able to determine the application on the basis of the given number of packets from any point in the flow [20].

Monitoring of the QoS can be done in a passive or an active mode. The passive mode relies mostly on time-based statistics, which are obtained directly from the flow passing the measurement point. This way, we can assess the jitter, the burstiness and the transmission speed (both download and upload). Unfortunately, it is not possible to receive the information about the packet loss or the delay for other than TCP streams while using this method. For that reason, additional tools performing active measurements must be involved in the process of estimating the QoS. One option is to use the ping-based approach, as it can measure both the delay and packet loss. Unfortunately, other issues can arise. Ping requests and responses are often blocked by network administrators, or their priority is modified (decreased to save the bandwidth or increased to cheat the users about the quality of the connection). Other options include sending IP packets with various TTL and awaiting *Time Exceeded* ICMP messages, which are usually allowed to be transmitted

in all the networks and their priority is not changed. Active measurements must be done in both directions (from the user and from the remote side). The total packet loss and the delay can be calculated as the sum of the delays and the packet losses from both directions of the flow. Furthermore, the knowledge of the direction that causes the problems can be used to assess if the problems are located in the local network or somewhere outside.

IX. CONCLUSION

This paper shows a novel method for assessing the Quality of Service in computer networks. Our approach involves a group of volunteers from the target network to participate in the initial training of the system, and later in the self-learning process. The accurate data obtained from the volunteers are used by the C5.0 MLA to create the per-application profiles of the network traffic as classification decision trees. The centralized measurement system uses the decision trees to determine the applications associated with the flows passing through the measurement point. This knowledge allows us to precisely define the QoS requirements for each particular flow. To assess the QoS level two methods are proposed: the passive and the active one.

X. ACKNOWLEDGMENT

This research work was conducted in the context of a PhD project in the Section for Networking and Security in the Department of Electronic Systems at Aalborg University. We are also very grateful, for the possibility to install and test our Volunteer-Based System, to *Gimnazjum nr 3 z Oddziałami Integracyjnymi i Dwujęzycznymi imienia Karola Wojtyły w Mysłowicach*, a high school in Poland.

REFERENCES

- [1] Tomasz Bujlow, Tahir Riaz, and Jens Myrup Pedersen. A method for Assessing Quality of Service in Broadband Networks. In *Proceedings of the 14th International Conference on Advanced Communication Technology (ICACT)*, pages 826–831. IEEE, Phoenix Park, PyeongChang, Korea, February 2012. Accessible: <http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=6174795>.
- [2] Gerhard Haßlinger. Implications of Traffic Characteristics on Quality of Service in Broadband Multi Service Networks. In *Proceedings of the 30th EUROMICRO Conference (EUROMICRO'04)*, pages 196–204. IEEE, Rennes, France, September 2004. DOI: [10.1109/EURMIC.2004.1333372](https://doi.org/10.1109/EURMIC.2004.1333372).
- [3] Thomas M. Chen and Lucia Hu. Internet Performance Monitoring. *Proceedings of the IEEE*, 90(9):1592–1603, September 2002. DOI: [10.1109/JPROC.2002.802006](https://doi.org/10.1109/JPROC.2002.802006).
- [4] LIRNEasia Broadband QoSE Benchmarking project, 2008. [Online]. Available: <http://lirneasia.net/projects/2008-2010/indicators-continued/broadband-benchmarking-qos-20/>.
- [5] K. v. M. Naidu, Debmalya Panigrahi, and Rajeev Rastogi. Detecting Anomalies Using End-to-End Path Measurements. In *Proceedings of the 27th Conference on Computer Communications IEEE INFOCOM 2008*, pages 16–20. IEEE, Phoenix, Arizona, USA, April 2008. DOI: [10.1109/INFOCOM.2008.248](https://doi.org/10.1109/INFOCOM.2008.248).
- [6] Aboagela Dogman, Reza Saatchi, and Samir Al-Khayatt. An Adaptive Statistical Sampling Technique for Computer Network Traffic. In *Proceedings of the 7th International Symposium on Communication Systems Networks and Digital Signal Processing (CSNDSP 2010)*, pages 479–483. IEEE, Newcastle upon Tyne, England, United Kingdom, July 2010.

- [7] Baek-Young Choi, Jaesung Park, and Zhi-Li Zhang. Adaptive Random Sampling for Traffic Load Measurement. In *Proceedings of the IEEE International Conference on Communications (ICC'03)*, volume 3, pages 1552–1556. IEEE, Anchorage, Alaska, USA, May 2003. DOI: [10.1109/ICC.2003.1203863](https://doi.org/10.1109/ICC.2003.1203863).
- [8] Shao Liu, Mung Chiang, Mathias Jourdain, and Jin Li. Congestion Location Detection: Methodology, Algorithm, and Performance. In *Proceedings of the 17th International Workshop on Quality of Service (IWQoS 2009)*, pages 1–9. IEEE, Charleston, South Carolina, USA, July 2009. DOI: [10.1109/IWQoS.2009.5201404](https://doi.org/10.1109/IWQoS.2009.5201404).
- [9] Jun Li, Shunyi Zhang, Yanqing Lu, and Junrong Yan. Real-time P2P traffic identification. In *Proceedings of the IEEE Global Telecommunications Conference (IEEE GLOBECOM 2008)*, pages 1–5. IEEE, New Orleans, Louisiana, USA, December 2008. DOI: [10.1109/GLOCOM.2008.ECP.475](https://doi.org/10.1109/GLOCOM.2008.ECP.475).
- [10] Kartheepan Balachandran and Jacob Honoré Broberg. Volunteer-Based Distributed Traffic Data Collection System. Master's thesis, Aalborg University, Department of Electronic Systems, Denmark, June 2010.
- [11] Riyad Alshammari and A. Nur Zincir-Heywood. Machine Learning based encrypted traffic classification: identifying SSH and Skype. In *Proceedings of the IEEE Symposium on Computational Intelligence for Security and Defense Applications (CISDA 2009)*, pages 1–8. IEEE, Ottawa, Ontario, Canada, July 2009. DOI: [10.1109/CISDA.2009.5356534](https://doi.org/10.1109/CISDA.2009.5356534).
- [12] Sven Ubik and Petr Žejdl. Evaluating application-layer classification using a Machine Learning technique over different high speed networks. In *Proceedings of the Fifth International Conference on Systems and Networks Communications (ICSNC)*, pages 387–391. IEEE, Nice, France, August 2010. DOI: [10.1109/ICSNC.2010.66](https://doi.org/10.1109/ICSNC.2010.66).
- [13] Ying Zhang, Hongbo Wang, and Shiduan Cheng. A Method for Real-Time Peer-to-Peer Traffic Classification Based on C4.5. In *Proceedings of the 12th IEEE International Conference on Communication Technology (ICCT)*, pages 1192–1195. IEEE, Nanjing, China, November 2010. DOI: [10.1109/ICCT.2010.5689126](https://doi.org/10.1109/ICCT.2010.5689126).
- [14] Jing Cai, Zhibin Zhang, and Xinbo Song. An analysis of UDP traffic classification. In *Proceedings of the 12th IEEE International Conference on Communication Technology (ICCT)*, pages 116–119. IEEE, Nanjing, China, November 2010. DOI: [10.1109/ICCT.2010.5689203](https://doi.org/10.1109/ICCT.2010.5689203).
- [15] Riyad Alshammari and A. Nur Zincir-Heywood. Unveiling Skype encrypted tunnels using GP. In *Proceedings of the 2010 IEEE Congress on Evolutionary Computation (CEC)*, pages 1–8. IEEE, Barcelona, Spain, July 2010. DOI: [10.1109/CEC.2010.5586288](https://doi.org/10.1109/CEC.2010.5586288).
- [16] Kartheepan Balachandran, Jacob Honoré Broberg, Kasper Revsbech, and Jens Myrup Pedersen. Volunteer-Based Distributed Traffic Data Collection System. In *Proceedings of the 12th International Conference on Advanced Communication Technology (ICACT 2010)*, volume 2, pages 1147–1152. IEEE, Phoenix Park, PyeongChang, Korea, February 2010.
- [17] Tomasz Bujlow, Kartheepan Balachandran, Tahir Riaz, and Jens Myrup Pedersen. Volunteer-Based System for classification of traffic in computer networks. In *Proceedings of the 19th Telecommunications Forum TELFOR 2011*, pages 210–213. IEEE, Belgrade, Serbia, November 2011. DOI: [10.1109/TELFOR.2011.6143528](https://doi.org/10.1109/TELFOR.2011.6143528).
- [18] Tomasz Bujlow, Kartheepan Balachandran, Sara Līgaard Nørgaard Hald, Tahir Riaz, and Jens Myrup Pedersen. Volunteer-Based System for research on the Internet traffic. *TELFOR Journal*, 4(1):2–7, September 2012. Accessible: <http://journal.telfor.rs/Published/Vol4No1/Vol4No1.aspx>.
- [19] Volunteer-Based System for Research on the Internet, 2012. [Online]. Available: <http://vbsi.sourceforge.net/>.
- [20] Tomasz Bujlow, Tahir Riaz, and Jens Myrup Pedersen. A method for classification of network traffic based on C5.0 Machine Learning Algorithm. In *Proceedings of ICNC'12: 2012 International Conference on Computing, Networking and Communications (ICNC): Workshop on Computing, Networking and Communications*, pages 244–248. IEEE, Maui, Hawaii, USA, February 2012. DOI: [10.1109/ICNC.2012.6167418](https://doi.org/10.1109/ICNC.2012.6167418).
- [21] Tomasz Bujlow, Tahir Riaz, and Jens Myrup Pedersen. Classification of HTTP traffic based on C5.0 Machine Learning Algorithm. In *Proceedings of the Fourth IEEE International Workshop on Performance Evaluation of Communications in Distributed Systems and Web-based Service Architectures (PEDISWESA 2012)*, pages 882–887. IEEE, Cappadocia, Turkey, July 2012. DOI: [10.1109/ISCC.2012.6249413](https://doi.org/10.1109/ISCC.2012.6249413).
- [22] Jason But, Philip Branch, and Tung Le. Rapid identification of BitTorrent Traffic. In *Proceedings of the 35th Annual IEEE Conference on Local Computer Networks (LCN)*, pages 536–543. IEEE, Denver, Colorado, USA, October 2010. DOI: [10.1109/LCN.2010.5735770](https://doi.org/10.1109/LCN.2010.5735770).
- [23] Li Jun, Zhang Shunyi, Lu Yanqing, and Zhang Zailong. Internet Traffic Classification Using Machine Learning. In *Proceedings of the Second International Conference on Communications and Networking in China (CHINACOM '07)*, pages 239–243. IEEE, Shanghai, China, August 2007. DOI: [10.1109/CHINACOM.2007.4469372](https://doi.org/10.1109/CHINACOM.2007.4469372).
- [24] Yongli Ma, Zongjue Qian, Guochu Shou, and Yihong Hu. Study of Information Network Traffic Identification Based on C4.5 Algorithm. In *Proceedings of the 4th International Conference on Wireless Communications, Networking and Mobile Computing (WiCOM '08)*, pages 1–5. IEEE, Dalian, China, October 2008. DOI: [10.1109/WiCom.2008.2678](https://doi.org/10.1109/WiCom.2008.2678).
- [25] Wei Li and Andrew W. Moore. A Machine Learning Approach for Efficient Traffic Classification. In *Proceedings of the Fifteenth IEEE International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunications Systems (MASCOTS'07)*, pages 310–317. IEEE, Istanbul, Turkey, October 2007. DOI: [10.1109/MASCOTS.2007.2](https://doi.org/10.1109/MASCOTS.2007.2).
- [26] Is See5/C5.0 Better Than C4.5, 2009. [Online]. Available: <http://www.rulequest.com/see5-comparison.html>.



Tomasz Bujlow is working as a Ph.D. Student in the Section for Networking and Security (NetSec) in the Department of Electronic Systems at Aalborg University in Denmark. He received his Master of Science in Computer Engineering from Silesian University of Technology in Poland in 2008, specializing in Databases, Computer Networks and Computer Systems. Previously, he obtained his Bachelor of Computer Engineering from University of Southern Denmark in 2009, specializing in software engineering and system integration. His research interests include methods for measurement of Quality of Service and traffic classification in computer networks. He is also a Cisco Certified Network Professional (CCNP) since 2010.



Sara Līgaard Nørgaard Hald is working as a Ph.D. student in the Section for Networking and Security (NetSec) in the Department of Electronic Systems at Aalborg University in Denmark. She received her Master of Science in Computer Engineering and Management from Aalborg University in 2002, and has since worked for the Danish Defense and as a consultant specializing in enterprise architecture and cybersecurity. Research interests include threat assessments and attack detection in dedicated networks.



Tahir Riaz is working as an Assistant Professor in the Section for Networking and Security (NetSec) in the Department of Electronic Systems at Aalborg University in Denmark. He received his Master and PhD degrees in Electrical and Electronics Engineering, specializing in Network Planning and Management, from Aalborg University in 2004 and 2008, respectively. He has also worked in Nokia, Linköping, Sweden. He has authored or co-authored over 70 papers published in conferences and journals. His research interests include access and backbone fiber optic networks, network planning and design, architecture of next generation radio of fiber networks, reliability and QoS issues in large-scale access and core network infrastructures, performance and optimization in networks.



Jens Myrup Pedersen is working as an Associate Professor and the head of the Section for Networking and Security (NetSec) in the Department of Electronic Systems at Aalborg University. His current research interests include network planning, traffic monitoring, and network security. He obtained his Master of Science in Mathematics and Computer Science from Aalborg University in 2002, and his PhD in Electrical Engineering also from Aalborg University in 2005. He is an author/co-author of more than 70 publications in international conferences and journals, and has participated in Danish, Nordic and European funded research projects. He is also a board member of a number of companies within technology and innovation.