

Towards QoS-aware load distribution in heterogeneous networks

Christian Niephaus^{*†}, Mathias Kretschmer^{*}, Gheorghita Ghinea[†]

^{*}Fraunhofer FOKUS, Sankt Augustin, Germany

[†]Brunel University, London, United Kingdom

christian.niephaus@fokus.fraunhofer.de, mathias.kretschmer@fokus.fraunhofer.de, george.ghinea@brunel.ac.uk

Abstract—Enabling broadband internet connectivity of 30 mbps and more is an ambiguous goal of the European Digital Agenda, particularly in rural and remote regions. Not relying on a single access technology but using multiple simultaneously is believed to be a promising option to meet this objective. However, simply using the available connections in parallel and distributing traffic arbitrarily among them despite their different characteristics might still lead to an unacceptable service quality due to the heterogeneity. Instead, methods that are sophisticated are required, which on one hand takes the Quality-of-Service (QoS) requirements of the various applications into account and on the other hand is aware of the different network characteristics. In this work, we discuss the main challenges which occur when utilizing multiple access technologies in parallel and we propose an architecture addressing those issues. Moreover, we present some preliminary validation results, which show the benefit of our approach.

Index Terms—QoS, load distribution, heterogeneous access networks, link abstraction, integration satellite and terrestrial networks

I. INTRODUCTION

Broadband Internet connectivity has become increasingly important over the recent years and is nowadays considered “a crucial factor to realize economic growth” [1], which enables the development of new services and applications, as described in the European Digital Agenda. Hence, in this agenda the European Commission sets the objective to enable broadband Internet connections of at least 30 mbps to be available to all EU citizens and 100 mbps to at least half of European households by 2020. Comparable institutions all over the world have set similar goals. While this might be realistic in densely populated, urban areas it is a challenging task in rural and other difficult-to-serve areas where a roll-out of traditional terrestrial current broadband technologies, such as X-Digital subscriber line (xDSL), Fiber to the X (FTTx) or even Long Term Evolution (LTE) are economically unfeasible for operators.

Next generations of fixed satellite systems are seen as one possible solution to address this issue [2]. Those systems, which are scheduled to be operational by 2020, might lead to Terabit/s satellite systems and will decrease the cost per bit tremendously, which makes them suitable for reaching rural areas [3]. However, satellite links will still introduce

high latency on connections making it difficult for users to perceive a high user Quality-of-Experience (QoE) when latency intolerant applications are being used. For example, Voice over IP (VoIP) and other interactive applications should be serviced with a latency of not more than 100ms and a jitter below 50ms, whereas a video streaming application can easily tolerate a latency around 1s [4]. Obviously, the first cannot be achieved with Geostationary Earth orbit (GEO) satellite systems, which are nowadays used to provide internet access to end-users via satellite, while the first cannot [5].

Thus, using solely new satellite systems to provide broadband Internet access in rural and remote areas would not solve the problem since the QoS demands, as expected in 2020 by the users, independent of their location cannot be met due to the high latency, i.e. a user in a remote area expects the same high quality service as a user living in an urban region. Instead, complementing multiple heterogeneous access networks, as depicted in Fig. 1, can be a potential solution to provide the required bandwidth to the end users also in rural and remote areas while still allowing for a high QoS. By simultaneously using multiple connections higher availability, reliability and eventually a better QoE can be provided. We argue that to achieve this, a novel architecture is required to seamlessly integrate the satellite links into terrestrial networks in order to exploit the advantages of both worlds, i.e. the high bandwidth of the satellite system and the low latency of terrestrial access technologies.

In this paper, we present the concept how the complementation of different terrestrial and satellite access networks can be exploited to provide broadband connectivity in rural areas while maintaining a high QoS level. We identify which new components and mechanisms are required and describe their functionalities. The remainder of this paper is structured as follows: First, the actual challenges are described and the related work is being discussed. Afterwards an architecture addressing the issues identified above is presented and preliminary validation results are shown. We conclude with a summary and the planned future work. It should be noted that this work focuses on describing the challenges that need to be solved and different pieces required to exploit the potential of satellite and terrestrial network integration. This paper does

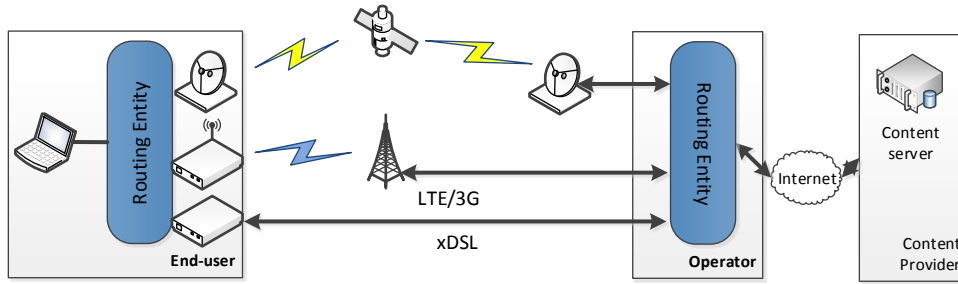


Fig. 1. Multiple connection scenario overview

not intent to provide a holistic solution to the problem.


II. PROBLEM STATEMENT

Simply distributing traffic equally on a per-packet basis among many connections, as it is often done in Local Area Networks (LANs) to increase the available bandwidth e.g. between two switches would be suboptimal if the connections are very different in their characteristics in terms of bandwidth, latency, jitter or loss. Connections might be over- or underloaded or applications need to be highly tolerant to jitter.


This is shown in Fig. 2, where the end-to-end latency and jitter values of a VoIP flow and a video streaming flow are shown, which are sent over an emulated combined terrestrial/satellite connection. The aggregated capacity of both connections together is 30mbps, the latency, however, is 10ms on the emulated terrestrial connection but 500ms on the satellite connection. The packets of both flows are distributed equally via both connections without considering any QoS requirements.

During the first approx. 10s only the voice flow (red curve) was active. It can be seen that during that time the latency is constantly 250ms, as shown in Figure 2(a), and the jitter is constantly 500ms, as depicted in Fig.2(b), which is obvious since two successive packets are not sent via the same path. Hence, they experienced either 500ms or 10ms latency leading to approximately 500ms jitter. Once the additional video traffic (green curve) has started the variation increases somewhat but the mean latency is still around 250ms and the jitter values are between 400 and 500ms. It is also obvious that the quality of the VoIP flow perceived by a potential user would be extremely bad and the service would be most likely not useable. Moreover, also the video flow quality would be extremely bad due to the high jitter [4]. It should also be noted that this kind of load distributing among the links will also cause extensive reordering which might tremendously decrease the TCP performance [6].

To improve the QoS and QoE, novel mechanisms are required to intelligently distribute the end-user's traffic among the available connections while taking the traffic QoS demands into account, in order to optimally use heterogeneous technologies concurrently. It is essential for a routing entity to know on one hand, the characteristics of the different connections and



(a) Latency



(b) Jitter

Fig. 2. Voice and video flow over a combined satellite and DSL link

on the other hand, the requirements of the active applications on the network. Referring to the aforementioned example, the VoIP call, which requires limited available bandwidth but a low latency and jitter should not be routed via a satellite link. In contrast to that, the video streaming traffic, which is more tolerant to latency, might require a high bandwidth but usually has limitations in terms of jitter, which makes it suitable to be routed via satellite link. Another, more complex example is online gaming. Usually before the actual game starts a large amount of data is transferred, followed by a period of only limited traffic during the actual game. In this first period, high bandwidth is required whereas during the second period low

latency is most crucial. In particular, the last example shows that deriving the QoS requirements of certain traffic is not a trivial task since traffic generated by one application might impose completely different demands on the network.

In order to cope with the heterogeneous technologies and their differences, technology agnostic *Link abstraction* or *Connectivity abstraction* is required so that different characteristics of each individual connection can be described in a systematic and efficient manner by a set of well-defined parameters, including capacity, typical latency and jitter and so forth. Due to integrated satellite links this link layer abstraction needs to deal with unidirectional links as well as multicast delivery mechanisms. This link abstraction have to be complemented by an abstract QoS description, which maps the QoS requirements of traffic onto a similar set of generic and abstract parameters. Along with a monitoring system that generates likewise abstract and technology independent events, this enables the routing entity to determine the best connection for each flow.

III. RELATED WORK

A. Traffic classification

In order to ensure a high QoS experience it is essential to identify the QoS requirements of the traffic and, thus, to classify the traffic. However, mainly due to the increasing capacity and availability of today's broadband networks, a wider range of services such as VoIP, peer-to-peer (P2P) applications for sharing audio and video files, etc. are used by today's customers. Compared to modem dial-up users this lead to a more complex behavior and traffic patterns.

The area of traffic classification has attract researchers enormously over the last year. Several authors proposed a taxonomy and a comparison of different traffic classification methods and techniques, e.g. [7], [8] or [9]. Usually all traffic classification techniques can be assigned to four main categories: First, a simple identification method based on the transport-layer ports used by the application. Second, a payload-based method (sometimes also referred to as Deep packet inspection (DPI)) which classifies traffic by analyzing the headers and the payload of packets. Third, a host-based approach, which identifies traffic by patterns of host behavior on different levels [10] and finally, classification methods which use machine learning techniques or pattern matching approaches to assign traffic to application types.

While the port-based method have been accurate some years ago, nowadays transport-layer port numbers can only provide limited information. As already described in e.g. [10], [11] or [12], several recent developments impact the accuracy of this method such as the re-use of well-known ports or protocols used to tunnel other protocols [11].

Obviously, the second, payload-based method, is the most accurate solution since it inspects and evaluates all headers and payload of the packet, assuming the traffic is not encrypted and the used protocols are known by the traffic classifier so that they can be interpreted [7]. Many reasons, however, render this method difficult to implement. Protocols, which

are using encryption or are proprietary, prevent effectively the decoding of (encapsulated) headers and payload. Moreover, in many countries privacy laws prohibits the inspection of packets above the network layer [8]. As also outlined in [9], the packet-based approaches scale poorly with increasing bandwidth and are quite resource-intensive since virtually a parser for every protocol that might occur on a link is required in order to allow an accurate classification. Particular since nowadays applications emerge more frequently, it might be much effort to keep the classifier up to date.

In contrast to the previous approaches host-based classification methods, for example the BLINC method [10], first associates hosts with the applications they are using and then identifies the traffic accordingly. [10] proposes to investigate the behavior of a host at three levels, namely a social level, a network level and an application level. At the social level, it is analyzed with how many other hosts it communicates with by examining source and destination IPs. The functional role of a host describes whether it is providing a service or is rather a consumer. This analysis can be done based on the source and destination port. Hosts which are using a single source port are usually providing a service on that port. Last, at the application level additional flow information are evaluated, such as the transport protocol, number of packets or bytes and so forth. In the case of BLINC, a library of host-based signatures is used to identify the concrete application. Such an approach requires a bidirectional flow in order to determine the traffic class, hence it can only be used reliably at the edges of a network.

Due to the aforementioned shortcomings of the Packet- and port-based approaches, new classification methods emerged which use the statistical characteristics of flows and data mining or machine learning algorithms for classifying traffic. Machine learning techniques work on so-called *features*, which are attributes of a flow, such as median inter-packet arrival time, mean packet length and so forth [13]. Finally, some approaches use a combination of these four method such as [14] or [15].

A major problem is encrypted and tunneled traffic, such as IPSec[16] encrypted VPN tunnels. In particular if the traffic classification is used for QoS provisioning, e.g. to prioritize VoIP traffic, it becomes important to identify the VoIP traffic among other encrypted Best Effort traffic. In [17] the authors have shown that identifying VoIP traffic in encrypted tunnels can increase significantly increase of performance if the network is heavily loaded.

Moreover, to the best knowledge of the authors these methods classify traffic into categories but do not derive the corresponding QoS parameters required to allow for proper traffic distribution on the available links.

B. Load Distribution

Also the field of Load Distribution is well-studied in both the academia and the industry world. [18] surveys Load distribution methods over multipath networks and identifies traffic splitting and path selection as the key components of

load distribution. The first component splits the traffic into units such as a packet, a flow or a flowlet and the latter determines which path is used for which unit of traffic. A main differentiation between concrete approaches is inter- and intra-flow parallelism. Particularly if Load distribution is used to increase the available bandwidth the traffic belonging to the same flow (or application) can be either send over multiple available connections in parallel (intra-flow) or only over a single one and another flow might use another parallel path (inter-flow). Since in the latter approach the path selection is performed for multiple, related packets, inter-flow mechanisms might help to prevent reordering, which avoids problems for TCP connections as mentioned above.

Load distribution itself can also occur on different layers in the Open Systems Interconnection (OSI) reference model. It is often used on the Network Layer and then called *Multipath routing* as this layer is responsible for routing but it can be also done on the Link Layer or the Transport layer. A special form of multipath routing and well known concept in routing and traffic engineering is Equal Cost Multipath (ECMP), which becomes important if multiple paths with the same link cost to the a destination exist, and the routing entity needs to decide, which of the available paths it uses [19].

According to [18] most of the path selection models can be categorized into four categories, namely Round-Robin, Packet-Info, Traffic-Condition-based and Network-Condition-based. The first distributes traffic units in a round robin manner on all available paths without taking any external information into account, or using information gained from the packet headers to distribute the traffic. Such a method has been used for the measurements presented in Section I. In contrast to that, Packet-Info-based models are using information from protocol headers to select the best path. Network-condition-based and Traffic-Condition-based models are considered *adaptive* models, which are able to react on dynamically changing conditions of the traffic or the networks, contrary to *non-adaptive* models such as the round robin models and some of the Packet-Info-based models. Traffic-Condition-based models are taking the traffic load and traffic characterization into account while Network-Condition-based models are focusing on the network conditions such as desired and actual load.

It should be noted that those four categories are not mutually exclusive and concrete implementations might combine two concepts, e.g. a protocol might take traffic and network conditions into account as well as other protocol information when selecting the best bath for a traffic unit. However, a concrete model can either be adaptive or non-adaptive.

RFC2991[21] also discusses multipath issues with respect to the unicast and multicast next-hop selection. Among others, the authors identify variable latencies and a variable path Maximum Transfer Unit (MTU) of different links as major concerns, which might arise when the traffic is distributed on a per-packet basis. Hence, [21] recommend to not split 'flows' among multiple paths, where the term flow is defined as the "granularity at which the router keeps state [...] for classes of traffic".

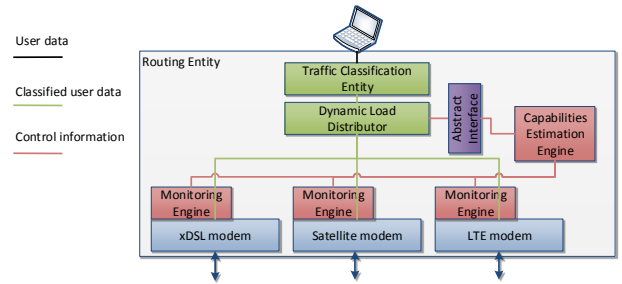


Fig. 3. Functional components overview

However, even the adaptive approaches do only take traffic characteristics such as packet size or size of flow into account but the methods are unaware concrete QoS requirements of different traffic flows and, thus, do not consider it when distributing the traffic, leading to situation where a e.g. an latency intolerant phone call might be routed via a high latency satellite link.

IV. PROPOSED OVERALL ARCHITECTURE

A network architecture has been defined to address those issues listed in section I. As depicted in Fig.1 the network consists of two sites, the end-user and the operator site, which are interconnected by multiple links of heterogeneous technologies and a routing entity on both sites, which selects the proper link. It should be noted that these connections not necessarily have to be a direct point-to-point link but rather might be a chain of links, e.g. a LTE connection might have one wireless hop from the LTE modem to the eNodeB and several, wireless or wired, hops from there to the operator's core network.

The core architecture is depicted in Fig.3. For the reason of space Fig.3 shows only the end user's edge of the network. However, the structure on the operator's edge is similar. The central component of the routing entity is the *Dynamic Load Distributor*, which selects a proper network to transmit a certain flow. It takes into account two input parameters, namely the QoS requirements of the traffic and the capabilities of the available connections. The first is provided by the *Traffic Classification Entity*, which identifies and classifies the traffic generated by an end user in order to detect the QoS requirements of it. The capabilities and current conditions of the available connections are provided by the *Capabilities Estimation Engine*. It implements a technology agnostic interface that provides an abstract description of the different connections. This interface is is exploited by the *Dynamic Load Distributor*, which can operate without knowing the technology specifics. Moreover, this abstraction layer also hides technology specific events and characteristics, such as a change in the Modulation and Coding Scheme (MCS) leading to a change in the available capacity on a link, and provides generic events to the *Dynamic Load Distributor*, e.g. *bandwidth changed to X* or *link down*.

Along with the Traffic Classification Entity, which provides the QoS information of the traffic in a similar abstract manner, the Dynamic Load Distributor can effectively select the most optimal connection for each flow.

In order to react also on other dynamic changes in the used networks the *Monitoring Engines* continuously monitor each connection as well as the traffic flows send over it. Since these Monitoring Engines might interact with the different modems and evaluate L1 and L2 information they are technology specific. However, the technology specific events are also hidden from the Dynamic Load Distributor by the abstraction layer.

Moreover, by exploiting the abstract description of both, QoS requirements and connection capabilities, the Dynamic Load Distributor is independent from methods, which are used to gain the necessary information. For example, in order to determine the capabilities of the available communication networks two general approaches can be chosen, either the networks can report their capabilities or the Capabilities Estimation Engine can (try to) identify these capabilities itself. Whereas the first needs an additional protocol to exchange these information the latter requires active measurements. Similarly, the communication requirements of applications can be determined. Either the application provides the necessary information to the Traffic Classification Engine or it has to identify the traffic and to derive the proper requirements itself.

V. VALIDATION

In order to show the benefit of the proposed architecture we have repeated the measurement presented in Section I with different configurations. First, we deactivated the satellite link and, hence, forced that all traffic was sent via the DSL link, which has a capacity of 10 mbps on the down-link. In the first phase of the test, when only the VoIP traffic has been sent the latency is within an acceptable region and the packet loss is zero, as depicted in Fig. 4. Once the additional video traffic has started the latency and the packet loss for both flows increased tremendously, as the DSL link was heavily overloaded. Hence, it would not be possible to use both services in parallel.

This changed when the high capacity satellite link has been activated. In order to take also the different QoS requirements of the traffic is taken into account, in contrast to the measurement present in Section I, this time the latency intolerant VoIP traffic was routed always via the DSL link whereas the video flow, which demands high bandwidth, was sent via the satellite link. The results are depicted in Fig.5. Admittedly, the video flow was still experiencing around 500 ms latency but only a negligible jitter, which is acceptable for a non-interactive video stream. Latency and jitter of the VoIP was obviously low as the the traffic is routed via the DSL link.

The traffic classification was statically configured and is classifying the traffic based on the destination ports. Moreover, also the Capabilities Estimation Engine was statically pre-configured to always report the overall bandwidth of the two links. It is easy to see that such a method is not suitable for real user traffic that is not artificially created. In a real deployment

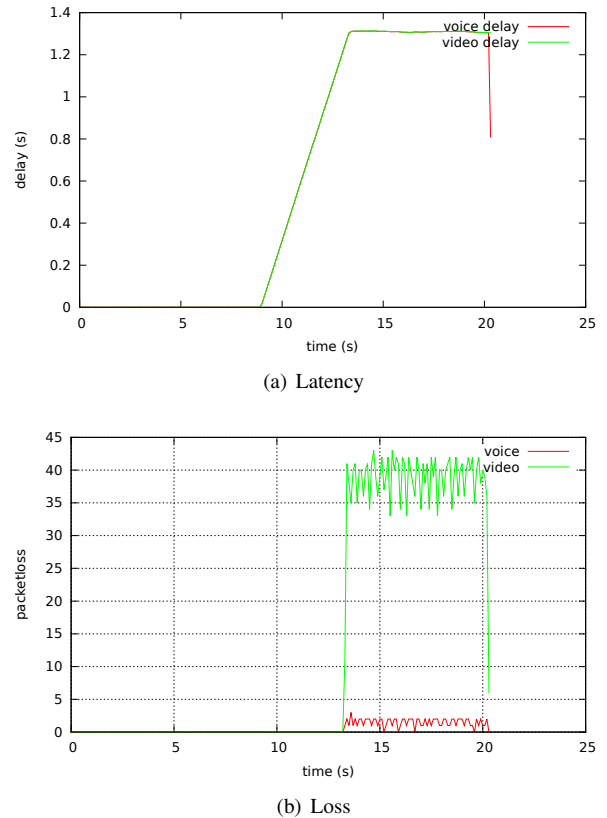


Fig. 4. Latency and Loss of a VoIP and a video flow over single DSL link

more matured techniques are required to deal with the complex mixture of typical end-user traffic, as mentioned in Section III.

VI. CONCLUSION AND FUTURE WORK

In this work, we argue that in order to accomplish the goal of the European Digital Agenda to provide at least 30 mbps to each European citizen next generation satellite systems will play an important role. Since they will heavily evolve during the next years leading to terabit/s satellite systems, costs per bit transmitted over a satellite link will decrease in the order of magnitude. However, to effectively exploit their advantages they need to be transparently integrated into existing terrestrial networks so that customers can use the best of both worlds, i.e. high capacity but also high latency satellite links as well as low latency but also low capacity terrestrial networks. Although satellite systems can provide high bandwidth connections in rural areas, high quality services, as demanded by today's customers, also require considering other QoS parameters besides the available bandwidth, such as latency, jitter or loss.

Moreover, we have shown that combining multiple access technologies in order to provide broadband connectivity in undeserved, remote and rural areas is not an easy and straightforward task as could have been expected, in particular if the considered technologies have a high degree of heterogeneity. It is essential that a routing entity does not solely select the proper network based on the destination IP address but rather takes the traffic QoS requirements as well as the capabilities




Fig. 5. Latency and jitter of a VoIP and a video flow over a combined satellite and DSL link

and the current status of the different available networks into account. It has become obvious that already existing solutions can only provide a partial solution for the problem.

Thus, we have presented a generic network architecture and identified the key building blocks needed to address this problem. We argue that a central Dynamic Load Distributor is required which shall work on an abstract level without knowing the technology specifics of each connection. Hence, we advocate for an abstract interface that provides the capabilities of each connection in an abstract and technology agnostic manner along with an abstract description of the QoS requirements of each flow. By using this architecture, the system can be tailored to cope with current and future access technologies.

Finally, we showed some preliminary results of measurements that show a clear benefit of the proposed architecture.

Future work will mainly focus on instantiating the proposed architecture, which includes defining and implementing its main components. The main work in this regard will be to define the Dynamic Load Distributor in order to ensure a scalable system able to cope with the increasing bandwidth in future systems.

ACKNOWLEDGMENT

This work has been supported by the BATS research project which is funded by the European Union Seventh Framework

Programme under contract n317533. The views and conclusions contained here are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the BATS project or the European Commission.

REFERENCES

- [1] EUROPEAN COMMISSION, "Digital agenda for europe scoreboard 2012," June 2012.
- [2] G. Peters, "Satellite delivery of next generation broadband access to the uk," in *Advanced satellite multimedia systems conference (asma) and the 11th signal processing for space communications workshop (spsc), 2010 5th*, pp. 141–146, 2010.
- [3] B. Evans, P. Thompson, L. Castanet, M. Bousquet, and T. Mathiopoulos, "Concepts and technologies for a terabit/s satellite," in *SPACOMM-2011*, (Budapest, Hungary), 2011.
- [4] I. T. Recommendation, "Y.1541: Network Performance Objectives for IP-Based Services," tech. rep., International Telecommunication Union, 2003.
- [5] X. Luan, J. Wu, X. Xu, S. Ren, and H. Xiang, "Research on the propagation delay characteristic of multi-beam geo satellite communications system," in *Advanced Communication Technology (ICACT), 2011 13th International Conference on*, pp. 620–623, 2011.
- [6] C. Arthur, A. Lehane, and D. Harle, "Keeping order: Determining the effect of tcp packet reordering," in *Networking and Services, 2007. ICNS. Third International Conference on*, pp. 116–116, 2007.
- [7] A. Callado, C. Kamienski, G. Szabo, B. Gero, J. Kelner, S. Fernandes, and D. Sadok, "A survey on internet traffic identification," *Communications Surveys Tutorials, IEEE*, vol. 11, pp. 37–52, quarter 2009.
- [8] A. Dainotti, A. Pescape, and K. Claffy, "Issues and future directions in traffic classification," *Network, IEEE*, vol. 26, pp. 35–40, january-february 2012.
- [9] H. Kim, k. claffy, M. Fomenkov, D. Barman, M. Faloutsos, and K. Lee, "Internet Traffic Classification Demystified: Myths, Caveats, and the Best Practices," in *ACM SIGCOMM CoNEXT*, (New York, NY), ACM SIGCOMM CoNEXT, Dec 2008.
- [10] T. Karagiannis, K. Papagiannaki, and M. Faloutsos, "Blink: multilevel traffic classification in the dark," in *Proceedings of the 2005 conference on Applications, technologies, architectures, and protocols for computer communications, SIGCOMM '05*, (New York, NY, USA), pp. 229–240, ACM, 2005.
- [11] A. W. Moore and K. Papagiannaki, "Toward the accurate identification of network applications," in *In PAM*, pp. 41–54, 2005.
- [12] L. Bernaille, R. Teixeira, I. Akodkenou, A. Soule, and K. Salamatian, "Traffic classification on the fly," *SIGCOMM Comput. Commun. Rev.*, vol. 36, pp. 23–26, Apr. 2006.
- [13] T. Nguyen and G. Armitage, "A survey of techniques for internet traffic classification using machine learning," *Communications Surveys Tutorials, IEEE*, vol. 10, pp. 56–76, quarter 2008.
- [14] G. Szabo, I. Szabo, and D. Orincsay, "Accurate traffic classification," in *World of Wireless, Mobile and Multimedia Networks, 2007. WoWMoM 2007. IEEE International Symposium on a*, pp. 1–8, june 2007.
- [15] G. Aceto, A. Dainotti, W. de Donato, and A. Pescape, "Portload: Taking the best of two worlds in traffic classification," in *INFOCOM IEEE Conference on Computer Communications Workshops, 2010*, pp. 1–5, march 2010.
- [16] S. Kent and R. Atkinson, "Security Architecture for the Internet Protocol." RFC 2401 (Proposed Standard), Nov. 1998. Obsoleted by RFC 4301, updated by RFC 3168.
- [17] T. Yildirim and P. Radcliffe, "Voip traffic classification in ipsec tunnels," in *Electronics and Information Engineering (ICEIE), 2010 International Conference On*, vol. 1, pp. V1–151–V1–157, aug. 2010.
- [18] S. Prabhavat, H. Nishiyama, N. Ansari, and N. Kato, "On load distribution over multipath networks," *Communications Surveys Tutorials, IEEE*, vol. 14, pp. 662–680, quarter 2012.
- [19] C. Hopps, "Analysis of an Equal-Cost Multi-Path Algorithm." RFC 2992 (Informational), Nov. 2000.
- [20] D. Thaler and C. Hopps, "Multipath Issues in Unicast and Multicast Next-Hop Selection." RFC 2991 (Informational), Nov. 2000.