

Waterfall Traffic Classification: A Quick Approach to Optimizing Cascade Classifiers

Paweł Foremski¹ · Christian Callegari² · Michele Pagano²

© The Author(s) 2016. This article is published with open access at Springerlink.com

Abstract Heterogeneous wireless communication networks, like 4G LTE, transport diverse kinds of IP traffic: voice, video, Internet data, and more. In order to effectively manage such networks, administrators need adequate tools, of which traffic classification is the basis for visualizing, shaping, and filtering the broad streams of IP packets observed nowadays. In this paper, we describe a modular, cascading traffic classification system—the Waterfall architecture—and we extensively describe a novel technique for its optimization—in terms of CPU time, number of errors, and percentage of unrecognized flows. We show how to significantly accelerate the process of exhaustive search for the best performing cascade. We employ five datasets of real Internet transmissions and seven traffic analysis methods to demonstrate that our proposal yields valid results and outperforms a greedy optimizer.

Keywords Network management · Convergent networks · Traffic classification · Machine learning

1 Introduction

Internet traffic *classification*—or *identification*—is the act of matching IP packets with the computer program or communication protocol that generated them [1]. It resembles an “Internet microscope”, which lets us to look at a given network link, see the traffic

✉ Paweł Foremski
pjf@iitis.pl
Christian Callegari
c.callegari@iet.unipi.it
Michele Pagano
m.pagano@iet.unipi.it

¹ The Institute of Theoretical and Applied Informatics of the Polish Academy of Sciences, Bałtycka 5, 44-100 Gliwice, Poland

² Department of Information Engineering, University of Pisa, Via Caruso 16, 56122 Pisa, Italy

flowing, and identify various types of IP flows. Another useful metaphor to (TC) is listening to two foreigners talking nearby and recognizing their human language. Quite often, we are able to identify an unfamiliar language or dialect even if we cannot fully understand it. Similarly, the TC problem is recognizing network protocols given their traffic, without interest in their full information content. Moreover, knowing the protocols behind IP flows makes networks easier to manage. For instance, TC is important for network monitoring: if we want to visualize the traffic flowing through a router, it is useful to know its components. TC also helps network security officers to reveal and track suspicious network activity. It is used for implementing (QoS) schemes, traffic shaping, and packet filtering. In convergent networks, TC is the mechanism that enables separate routing policies for voice, video, and data traffic.

A single IP packet alone is difficult to classify, as there is no application name in the packet headers. In the past, the service port number was used for discriminating the traffic class [2], but this became ineffective due to the raise of Peer-to-Peer (P2P) traffic in the early 2000s [3]. A popular and de facto standard method used nowadays is Deep Packet Inspection (DPI): pattern matching on full packet contents [4]. However, although being more accurate than port-based classification, it requires more computing power and brings privacy concerns. Moreover, pervasive encryption and other issues make DPI increasingly irrelevant [5, 6]. Instead, modern classifiers investigate groups of packets to find distinguishing features of specific application, rather than of single packets. Usually, a flow of packets is statistically summarized—for example, using the average packet size and inter-packet arrival time—and the resultant *feature vector* is classified using a Machine Learning (ML) algorithm [7]. Such methods are more reliable: the overall behavior of a particular protocol or host is examined instead of seeking for a strict match in a few packets.

The current challenge in TC is that in future it will have to deal with an increasing adoption of encryption, encapsulation, multi-channel techniques, and with the tremendous growth of the Internet [8]. Inevitably, the TC problem is becoming a very complex task that needs breaking into subproblems to keep it tractable. Recent papers proposed various interesting techniques tailored at subproblems in TC [9–11], but so far few authors addressed the problem of combining these proposals to work together. Thus, in this paper, we describe our method for integrating different traffic classifiers—the Waterfall architecture [12]—and we introduce a novel algorithm for optimizing such systems.

In more detail, we will show how to apply a Multiple Classifier Systems (MCS) technique called *cascade classification* [13] to build a modular TC system optimized for a given computer network. The Waterfall architecture lets for dedicated classifiers for different types of network traffic, thus we believe our contribution is important for convergent networks. For example, (LTE) networks allow transmitting voice calls directly over IP, along with ordinary Internet data, which is known as Voice over LTE (VoLTE). Usually, the network will use a finite set of destination IP addresses for the VoLTE traffic. If one wants to identify IP traffic in such a network, Waterfall would allow separate classifiers for VoLTE traffic—which is simple, e.g. using the IP address—and for typical Internet data—which is more challenging. In overall, the system would effectively use computing resources by applying adequate methods to various services present in the heterogeneous network. Moreover, our optimization technique would further tune the system for desired goal, e.g. real-time traffic visualization. Comparing to our introductory work [14], the contribution of this paper is as follows:

1. We give an extended description of our novel method for optimizing classification cascades (Sects. 2 and 3).

2. We describe how to implement our algorithm recursively, and we reflect on its time complexity (Sect. 3.3).
3. We extensively validate our proposal on a new dataset with reliable ground-truth information, and on 7 classification modules total (Sect. 4).
4. We compare our proposal with myopic optimizer (Sect. 4.4).
5. We release an open source implementation of our proposal as a publicly available module (Sect. 5).

We begin our paper with Sects. 2 and 3 describing the Waterfall architecture and our optimization method, respectively. Then, we present the experimental results in Sect. 4. We conclude our work in Sect. 5.

2 Cascade Traffic Classification

The field of network traffic classification needs a method for integrating results of various research activities. Many papers in this area describe classification methods that in principle propose a set of traffic features tailored at a set of network protocols [1, 9–11, 15–17]. Researchers promote their methods for classifying network traffic, which are usually quite effective, but none of them is able to exploit all observable phenomena in the Internet traffic and identify all kinds of protocols.

The question arises: could we integrate these approaches into one system, so that we move forward, building on the achievements of our colleagues? How would this improve classification systems, in terms of accuracy, functionality, completeness, and speed? Answering these questions can open new perspectives for traffic classification. A robust method for combining classifiers can promote research that is more focused on new phenomena in the Internet, rather than addressing the same old issues.

In this Section we describe Waterfall: a modular architecture for traffic identification systems, which we introduced in [12]. Waterfall allows existing classification methods to complement each other, which makes the system as a whole capable of providing higher performance than could be achieved by any of the constituent modules.

2.1 Background

A naïve approach to the integration problem would be to survey recent papers for traffic features and use them as long feature vectors, classified with a decent machine learning algorithm. Even with adequate techniques employed, this could quickly lead us to the *curse of dimensionality* [18]: an exponential growth in the demand for training data as the feature space dimensionality increases. Besides, network flows differ in the set of available features, e.g. only a part of Internet flows evoke DNS queries [10]. Some features need more packets to be computed: e.g. port number is available after one packet, whereas payload statistics need several tens of packets [11]. This means that different tools are needed for different protocols: some flows can be classified immediately using simple methods, while others need more sophisticated analysis. Finally, from the software engineering point of view, a big, monolithic system could be difficult to develop and maintain.

Instead, researchers adopt multi-classification—in particular the Behavior Knowledge Space (BKS) combination method that fuses the outputs of many classifiers into one final decision. In principle, the idea behind BKS is to ask all classifiers for their answers on a particular problem \mathbf{x} and then query a look-up table \mathbf{T} for the the final decision. The

table **T** is constructed during training of the system, by learning the behavior of classifiers on a labeled dataset. For example, if an ensemble of 3 classifiers replies (*A*, *B*, *A*) for a sample with a ground-truth label of *B*, then the cell in **T** under index (*A*, *B*, *A*) is *B* (see [13], p. 128). This powerful technique can increase the performance of TC systems—as shown by Dainotti et al. [19]—but comparing to Waterfall, it inherently requires *all* modules to be run on each flow, with the drawback that the more modules are used, the more processing power is required.

2.2 The Waterfall Architecture

Waterfall applies the idea of multi-classification, but queries the constituent classifiers in sequential manner instead of parallel. It employs *cascade classification*, of which Kuncheva writes in her book on multi-classification: “cascade classifiers seem to be relatively neglected although they could be of primary importance for real-life applications.” (in [13], p. 106). We argue that cascade classification is a powerful and effective technique for combining algorithms that identify Internet traffic.

The Waterfall idea is presented in Fig. 1. The input to the system is an IP flow—a feature vector **x**—which contains all the features required by all modules, but a particular module will usually use only a subset of **x**.

The system sequentially evaluates *selection criteria* that decide which *classification modules* to use for the problem **x**. If a particular criterion is fulfilled, the associated module is run. If it succeeds, the algorithm finishes. Otherwise, or if the criterion was not satisfied, the process advances to the next step. When there are no more modules to try, the flow gets rejected and is labeled as “Unknown”. More precisely,

$$Dec_i(\mathbf{x}) = \begin{cases} Class_i(\mathbf{x}) & Crit_i(\mathbf{x})\text{ satisfied} \wedge Class_i(\mathbf{x})\text{ successful} \\ Dec_{i+1}(\mathbf{x}) & \text{otherwise} \end{cases} \tag{1}$$

$$Dec_{n+1}(\mathbf{x}) = \text{Reject} \tag{2}$$

where Dec_i is the decision taken at step $i = \{1, 2, \dots, n\}$, n is the number of modules, $Class_i(\mathbf{x})$ is the protocol identified by the module i , and $Crit_i(\mathbf{x})$ is the associated criterion.

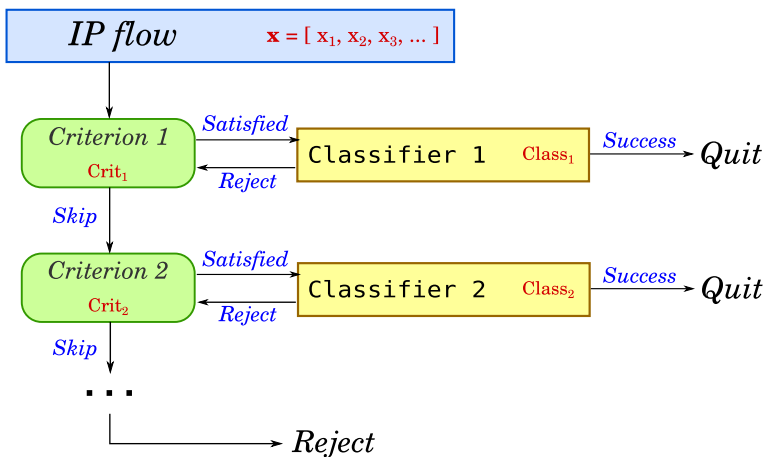


Fig. 1 The Waterfall architecture. A flow enters the system and is sequentially examined by the modules. In case of no successful classification, it is rejected

The selection criteria are designed to skip ineligible classifiers quickly. For example, in order to implement a module that identifies traffic by analyzing the packet payload sizes, the criterion could check if at least 5 packets with payload data were already sent in each direction. Only if this condition is true, a machine learning algorithm is run to identify the protocol. However, probably a large amount of flows will be skipped, saving computing resources and avoiding classification with an inadequate method. On the other hand, if a flow satisfies this criterion, it will be analyzed with a method that does not need to support corner cases (that is, number of payload packets less than 5). The selection criteria are optional, i.e. if a module does not have an associated criterion, the classification is always run.

3 Waterfall Optimization

Now we will consider the problem of optimal cascade structure. Let F be a set of IP flows, and E be a set of n classification modules,

$$E = \{1, \dots, n\} \quad (3)$$

that we want to use for cascade classification of flows in F in an optimal way. In other words, we need to find a sequence of modules X ,

$$X = (x_1, \dots, x_m) \quad m \leq n, \quad x_i \in E, \quad x_i \neq x_j \quad \text{for } i \neq j \quad (4)$$

that minimizes a cost function C ,

$$C(X) = f(T_X) + g(E_X) + h(U_X) \quad (5)$$

where the terms T_X , E_X , and U_X respectively represent the total amount of CPU time used, the number of errors made, and the number of flows left unlabeled while classifying F with X . The terms f , g , and h denote arbitrary real-valued functions. Because $m \leq n$, some modules may be skipped in the optimal cascade. Note that U_X does not depend on the order of modules, because unrecognized flows always traverse till the end of the cascade.

3.1 Background

Cascade classification is a multi-classifier system implementing the classifier selection idea [13]. Interestingly, although first introduced in 1998 by Alpaydin and Kaynak [20], so far few authors considered the puzzle of optimal cascade configuration that would match our problem. In a 2006 paper, Chellapilla et al. [21] propose a cascade optimization algorithm that updates the rejection thresholds of the constituent classifiers. The authors apply an optimized depth first search to find the cascade that satisfies given constraints on time and accuracy. However, comparing with our work, the system does not optimize the module order. In another paper on this topic, published in 2008 by Abdelazeem [22], the author proposes a greedy approach for building cascades: start with a generic solution and sequentially prepend a module that reduces CPU time. Comparing with our work, the approach does not evaluate all possible cascade configurations and thus can lead to sub-optimal results. We will demonstrate this in Sect. 4 for an exemplary myopic optimizer.

Thus, we propose a new solution to the cascade classification problem, which is better suited for traffic classification than existing methods. Note that comparing with [21] we do not consider rejection thresholds as input values to the optimization problem. Instead, in

case of classifiers with tunable parameters, one could consider the same module parametrized with different values as separate modules, and apply our technique as well. For instance, a Bayes classifier with rejection thresholds on the posterior probability of 0.5, 0.75, 0.90 would be considered as three separate modules.

3.2 Proposed Solution

To find the optimal cascade, we propose to approximate the performance of every possible X by calculating the performance of each module on the entire dataset and then smartly combining the results. Note that for an accurate solution one would basically need to run the full classification process for all permutations of all combinations in E . This would take S experiments, where

$$S = \sum_{i=1}^n \frac{n!}{(n-i)!} \approx e \cdot n! \tag{6}$$

which is impractical even for small n . On another hand, fully theoretical models of the cost function seem infeasible too, due to the complex nature of the cascade and module interdependencies.

Thus, we propose a heuristic solution to the cascade optimization problem. The algorithm has two evaluation stages:

1. Static: classify all flows in F using each module in E , and
2. Dynamic: find the X sequence that minimizes $C(X)$.

3.2.1 Static Evaluation

In every step of stage A, we classify all flows in F using each single module $x \in E$. We measure the average CPU time used for flow selection and classification: $t_s^{(x)}$ and $t_c^{(x)}$. We store each output flow identifier in one of the three outcome sets, depending on the result: $F_S^{(x)}$, $F_O^{(x)}$, or $F_E^{(x)}$. These sets hold respectively the flows that were skipped, properly classified, and improperly classified. Let us also introduce $F_R^{(x)}$,

$$F_R^{(x)} = F \setminus (F_S^{(x)} \cup F_O^{(x)} \cup F_E^{(x)}) \tag{7}$$

that is, the set of rejected flows. See Fig. 2 for an illustration of the module measurement procedure. As the result of every step, the performance of module x on F is fully characterized by a tuple $P^{(x)}$,

$$P^{(x)} = (F, t_s^{(x)}, t_c^{(x)}, F_S^{(x)}, F_O^{(x)}, F_E^{(x)}) \tag{8}$$

Finally, after n steps of stage A, we obtain n tuples: a model of our classification system, which is the input to stage B.

3.2.2 Dynamic Evaluation

Having all of the required experimental data, we can quickly estimate $C(X)$ for arbitrary X . Because f, g, h , are used only for adjusting the cost function—and can be modified by the




Fig. 2 Measuring performance of module $x \in E$

network administrator according to her needs (see Sect. 4.2)—we focus only on their arguments, i.e. the cost factors $T_X, E_X,$ and U_X .

Let $X = (x_1, \dots, x_i, \dots, x_m)$ represent certain order and choice of modules, and G_i represent the set of flows entering the module number i ,

$$G_1 = F \tag{9}$$

$$G_{i+1} = G_i \setminus (F_O^{(x_i)} \cup F_E^{(x_i)}) \quad 1 \leq i \leq m \tag{10}$$

then we estimate the cost factors using the following procedure:

$$T_X \approx \sum_{i=1}^m |G_i| \cdot t_s^{(x_i)} + |G_i \setminus F_S^{(x_i)}| \cdot t_c^{(x_i)} \tag{11}$$

$$E_X = \sum_{i=1}^m |G_i \cap F_E^{(x_i)}| \tag{12}$$

$$U_X = |G_{m+1}| \tag{13}$$

where $|G|$ denotes the number of flows in set G .

Note that the difference operator in Eq. 10 connects the static cost factors with the dynamic effects of a cascade. In stage A, our algorithm evaluates static performance of every module on the entire dataset F , but in stage B we want to simulate cascade operation, so we need to remove the flows that were classified in the previous steps. Thus, the operation in Eq. 10 is crucial.

Module performance depends on its position in the cascade, because preceding modules alter the distribution of traffic classes in the flows conveyed onward. For example, we can improve accuracy of a port-based classifier by putting a module designed for P2P in front of it, which should handle the flows that misuse the traditional port assignments.

3.3 Discussion

In our solution, instead of $e \cdot n!$ experiments (see Eq. 6), we simplified the optimization problem to n experiments and several computations, which in overall is much faster. Note that in case of adding a new module x_j to an already simulated cascade X , we can re-use previous computations:

$$G_j = U_X \quad (14)$$

$$T_{X+x_j} \approx T_X + |G_j| \cdot t_s^{(x_j)} + |G_j \setminus F_S^{(x_j)}| \cdot t_c^{(x_j)} \quad (15)$$

$$E_{X+x_j} = E_X + |G_j \cap F_E^{(x_j)}| \quad (16)$$

$$U_{X+x_j} = G_j \setminus (F_O^{(x_j)} \cup F_E^{(x_j)}) \quad (17)$$

Thus, we suggest searching for the minimum $C(X)$ in a recursive algorithm. However, although simulation is orders of magnitude faster than experimentation, we still check every possible cascade. This makes the time complexity of our algorithm factorial, considering set computations as the elementary operations. This might leave space for further improvements by the introduction of heuristics, possibly tuned to a specific cost function.

Moreover, note that the results depend on F : the optimal cascade depends on the protocols present in the traffic, and on the ground-truth labels. The presented method cannot provide the ultimate solution that would match every network, but it can optimize a specific cascade system for a specific network. We further discuss this issue in Sect. 4.

We assume that the flows are independent of each other, i.e. labeling a particular flow does not require information on any other flow. If such information is needed, e.g. flow DNS names, it should be extracted before the classification process starts. Thus, traffic analysis and flow classification must be separated to uphold this assumption. We successfully implemented such systems for our DNS-CLASS [10] and MUTRICS [12] classifiers.

In the next Section, we experimentally validate our method and show that it perfectly predicts E_X and U_X , and approximates T_X properly. The simulated cost follows the real cost, so we claim our proposal is valid and can be used in practice. We also analyze the trade-offs between speed, accuracy, and ratio of unlabeled flows, to stress out that the final choice of the cost function should depend on the purpose of the system.

4 Experimental Validation

Below we present the outcome of using real traffic datasets for experimental evaluation of our proposal. We ran four experiments:

1. comparing simulation with reality, which proves validity of Eqs. 11–13;
2. analyzing the effect of cost function parameters on the result, which demonstrates optimization for different goals;
3. optimizing on one dataset and using the cascade on another dataset, which evaluates stability;
4. comparing our optimization method with myopic optimization, which shows that our work is meaningful.

For the experiments, in general we used 5 datasets, summarized in Table 1. Datasets ASNET1 and ASNET2 were collected at the same ISP serving <500 domestic users, with an 8-month time gap. Dataset IITIS1 was collected at an academic network serving <50 researchers, at the same time as ASNET1. Dataset UNIBS1 was also collected at an academic network (University of Brescia,¹) but a few years earlier and using a reliable ground-truth

¹ Downloaded from <http://www.ing.unibs.it/ntw/tools/traces/>.

Table 1 Datasets used for experimental validation

Dataset	Start	Duration	Src. IP	Dst. IP (K)	Packets (M)	Bytes (G)	Avg. util (Mbps)	Avg. flows (/5 min)	Payload
<i>Asnet1</i>	2012-05-26	216 h	1800 K	1500	2500	1600	18	7.7 K	92 B
<i>Asnet2</i>	2013-01-24	168 h	2500 K	2800	2800	1800	26	12 K	84 B
<i>IITiSI</i>	2012-05-26	216 h	32 K	46	150	95	1.0	750	180 B
<i>Unibsl</i>	2009-09-30	58 h	27	1	30	26	0.9	110	0 B
<i>UPC1</i>	2013-02-25	65 days	90 K	18	37	33	51	68	Full
	2013-11-18	35 days	7.5 K	54	43	31	88	49	Full

information [23] (this dataset was anonymized). Finally, the UPC1 dataset was artificially generated—with manual simulation of different human behaviors—hence it contains full packet payloads and the names of applications that generated the traffic flows [24–26].

For the first 3 datasets, we established ground-truth using light DPI [27]. For UNIBS1 and UPC1, we used the supplied ground-truth information, which sometimes was challenging: for example, a *skype* process generates some HTTP traffic apart of the Skype protocol. For each dataset, we trained the modules using 60 % random sample of all flows, and used the remainder for testing. We considered only the first 10 s of each flow to resemble a near-immediate traffic identification.

Finally, in total we evaluated 7 classification modules, summarized in Table 2 [10, 12]. As additional traffic features, we used the transport protocol and destination port number for each module. Although we consider port numbers as an unreliable feature, they still can provide valuable hint for more sophisticated classification mechanisms. Note that the modules support the *reject option*, so each module can drop any flow if its not certain about the outcome.

4.1 Experiment 1

In the first experiment, we compare simulated cost factors with real values for arbitrary cascade configurations. We randomly selected 100,000 flows from each of the first 4 datasets and ran static evaluation on them. Next, we generated 100 random cascades, and for each cascade we ran both real and simulated classification. As a result, we obtained corresponding pairs of real and estimated values of T_X , E_X , and U_X .

The results for T_X are presented in Fig. 3. For E_X and U_X we did not observe a single error, i.e. our method perfectly predicted the real values. For CPU time estimations, we see a high correlation of 0.95, with little under-estimation of the real value. For all datasets, the estimation error was below 20 % for majority of evaluated cascades (with respect to the real value). The error was above 50 % only for 5 % of evaluated cascades.

We conclude that in general our method properly estimates the cost factors and we can use it to simulate different cascade configurations. Note that accurate prediction of the CPU time is not necessary for optimization: it is enough for the simulated time to be roughly proportional to the real value. Moreover, even the real values will vary depending e.g. on the CPU load due to other tasks executed in the background, which is difficult to predict.

Table 2 Waterfall modules used for experimental validation

Module	ML algorithm	Traffic features
dnsclass	Linear SVM	DNS name
dstip	Lookup table	Destination IP address
npkts	Random forest	Payload sizes: first 4 packets in+out
port	Lookup table	Destination port number
portsize	Lookup table	Payload sizes: first packet in+out
portname	Lookup table	DNS name
stats	Random forest	4 basic statistics of packet sizes and inter-arrival times

4.2 Experiment 2

In our second experiment we show the effect of tuning the system for 3 different goals: (a) minimizing the computation time, (b) minimizing errors, and (c) labeling as many flows as possible. We chose the following cost function:

$$C(X) = f(T_X) + g(E_X) + h(U_X) = (T_X)^a + (E_X)^b + (U_X)^c \quad (18)$$

with the default values of a , b , c equal to 0.95, 1.75, 1.20, respectively. We separately varied these values in range of 0–10, and observed the performance of the resultant cascades. For the sake of brevity, we ran the experiment for datasets ASNET1, ASNET2, and IITIS1 and for modules dnsclass, dstip, npkts, port, and portsize.

In Fig. 4, we present the results: dependence of cascade performance and module count on the cost function parameters. As expected, higher a exponent leads to faster classification and usually less errors, but with fewer modules in the cascade, and more unclassified flows as a consequence. Optimizing for accuracy—higher b exponent—leads to reduction of errors at the cost of higher number of flows left without a label. Finally, if we choose to classify as much traffic as possible (increasing the c exponent), the system will use all available modules, at the cost of higher CPU time and error rate.

In more detail, for time optimization, the *optimal* cascades are: port for ASNET1, portsize for ASNET2, and dnsclass for IITIS1. In the last case, dnsclass is preferred due to high percentage of DNS traffic in IITIS1. Instead, in case of accuracy optimization, the *optimal* cascades are: portsize, dnsclass, npkts, port for ASNET1, dstip, dnsclass, portsize for ASNET2, and dnsclass, port, dstip, portsize, npkts for IITIS1. Finally, optimizing for minimum percentage of unrecognized flows yields a common result for all datasets: dnsclass, dstip, npkts, port, portsize.

Note that the results depend on the cost function. We used a power function for presentation purposes, in order to easily show contrasting scenarios by small adjustments to the exponents. For specific purposes, a multi-linear function may be more appropriate, as it is often found in the literature, e.g. linear scalarization of multi-objective optimization problems. Moreover, more complex expressions—including thresholds on some parameters—can be used to find a classification system capable of real-time operation: given an expected amount of flows per second, one could find a cascade that is fast enough to handle the traffic while keeping the other cost factors at possible minimum.

We conclude that our proposal works and is adaptable, i.e. by varying the parameters we optimized the classification system for different goals.




Fig. 3 Experiment 1. Estimated classification time versus real classification time. *Dashed line* shows least-squares approximation. The Pearson product-moment correlation is 0.95

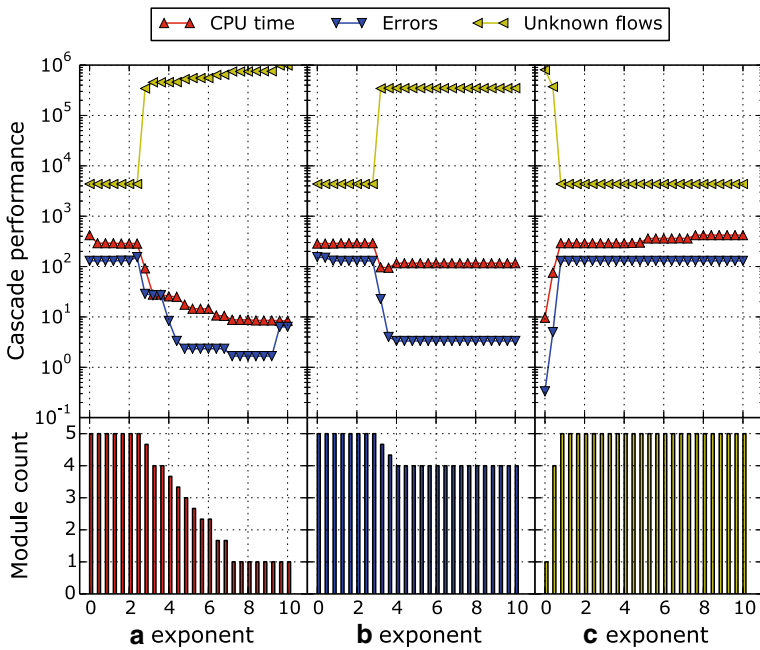


Fig. 4 Experiment 2. Optimizing the cascade for different goals: best classification time (*a* exponent), minimal number of errors (*b* exponent), and the lowest number of unlabeled flows (*c* exponent): the *plot* shows the averages for 3 datasets

4.3 Experiment 3

In the third experiment, we wanted to verify if the result of optimization is stable in time and space, i.e. if the optimal cascade stays optimal with time and changes of the network. We ran our optimization procedure for 4 datasets, obtaining different cascade configuration for each dataset. Next, we evaluated these configurations on all datasets and measured the increase in the cost function $C(X)$ compared with the original value. Note that we did not use the UNIBS1 dataset for this experiment, as it lacks packet payloads and hence needs different set of available modules.

Table 3 presents the results. We see that our proposal yielded results that are stable in time for the same network: the cascades found for ASNET1 and ASNET2, which are 8 months apart, are similar and can be exchanged with little decrease in performance. However, the cascades found for ASNET1 and ASNET2 gave 5–7 % worse performance compared with IITiS1, and 23–49 % worse performance on UPC1. We observed extreme decrease in performance when we varied both the network and time, especially when classifying UPC1 with cascade optimized for IITiS1.

We conclude that cascade optimization is specific to the network, but on the other hand our results suggest that an optimal cascade does not change significantly with time for given network. Thus, the network administrator does not need to repeat the optimization procedure frequently.

4.4 Experiment 4

In the last experiment, we compared our proposal with a greedy optimizer, i.e. a situation in which we select all modules in order of increasing CPU time. This resembles the basic approach in the original paper on Waterfall [12]: start with generic, heavy classifier, and prepend faster modules in front of it (see section 5 in [12]). Thus, for each module, we calculated the sum of t_s and t_c for each dataset separately, and ordered the modules from the fastest to the slowest. We used the results as cascade configurations, i.e. Waterfall systems configured with a conservative algorithm: “myopic” optimization.

On the other hand, we also optimized the system using our proposal, with the cost function given in Eq. 18, for a , b , c equal to 3.00, 1.75, 1.50, respectively. We chose these exponent values arbitrarily to show an example of time optimization: note that the a exponent (influencing the time cost factor) is the highest. Then, we used the results as cascade configurations, but optimized with an “optimal” algorithm.

Table 4 compares the results: in every case, our algorithm optimized the classification system to work faster and with less errors, usually with the same amount of unclassified flows. This demonstrates the point of cascade optimization: it brings performance

Table 3 Experiment 3. Result stability: relative increase in the cost $C(X)$, depending on the reference dataset used for determining the optimal cascade

Reference	Test dataset			
	<i>Asnet1</i> (%)	<i>Asnet2</i> (%)	<i>IITiS1</i> (%)	<i>UPC1</i> (%)
<i>Asnet1</i>		1.01	5.31	48.96
<i>Asnet2</i>	2.67		7.29	23.34
<i>IITiS1</i>	33.37	34.19		192.91
<i>UPC1</i>	14.51	11.11	31.77	

Table 4 Experiment 4. Average improvements compared to myopic cascade optimization

Dataset	Algorithm	Cascade configuration	Time (s)	Errors	Unknowns
Asnet1	Myopic	Portname, portsize, port, dstip, dnsclass, stats, npkts	89	40	886
	Optimal	Portsize, portname, dstip, dnsclass, npkts, port, stats	87	30	886
			+2 %	+26 %	0 %
Asnet2	Myopic	Portname, portsize, port, dstip, dnsclass, stats, npkts	141	49	817
	Optimal	Portsize, portname, dstip, dnsclass, npkts, port	139	22	1224
			+2 %	+55 %	-50 %
IITiS1	Myopic	Dnsclass, port, portname, portsize, dstip, stats, npkts	5.7	2.4	80
	Optimal	Port, portsize, npkts, stats	5.1	2.4	80
			+11 %	+0 %	0 %
Unibs1	Myopic	Portsize, port, dstip, stats, npkts	102	2017	13,892
	Optimal	Dstip, portsize, port, npkts, stats	91	1985	13,892
			+10 %	+2 %	0 %
UPC1	Myopic	Portname, port, portsize, dstip, dnsclass, stats, npkts	110	686	1746
	Optimal	Port, portname, dstip, portsize, dnsclass, npkts, stats	92	604	1746
			+16 %	+12 %	0 %
Average improvement			+8 %	+19 %	-10 %

improvements. Recall that UNIBS1 lacks packet payloads, hence we used 5 modules in general for this dataset instead of 7.

On average, the system worked 8 % faster compared with myopic time optimization, and reduced the error rate by 19 %. For ASNET2, it also resulted in higher number of unrecognized flows, but the increase is insignificant given the dataset size, and this cost factor was not the goal of optimization. For instance, if one wants a real-time traffic visualization system, then some small portion of flows might remain unrecognized without negative effect on the whole system. Thus, we conclude that our work is meaningful and can help network administrators to tune cascade TC systems better than ad-hoc tools.

5 Conclusions

We showed that our Waterfall architecture, together with the new optimization technique, lets for effective combining of traffic classifiers. We presented background on cascade classification (a multi-classifier variant) and employed it for identifying IP transmissions. Waterfall brakes the complex TC problem into smaller, independent modules, which are easier to manage. Moreover, we presented an optimization technique that automatically selects the set of best modules from a pool of available methods, and puts them in right order for maximized performance. By means of experimental validation we demonstrated

that our proposal works and can bring significant improvements to classification speed, accuracy, and number of recognized flows.

Our approach to optimizing Waterfall systems brings major improvements over ad-hoc methods. First, it reduces the time needed for optimization by orders of magnitude, by replacing experimentation on different cascades with simulation, which is much faster. Second, by performing an exhaustive search for the best solution, it finds better cascades than a greedy algorithm. However, due to the complex nature of the problem, it still requires a considerable amount of computations to check for all possible cascade configurations, which in practice limits the maximum size of the module pool.

We believe our contribution is important for managing convergent networks like LTE. Finally, in order to support further research in this area, we release an open source implementation of our proposal as an extension to the MUTRICS classifier, available at <https://github.com/iitis/mutrics>.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Foremski, P. (2013). On different ways to classify Internet traffic: A short review of selected publications. *Theoretical and Applied Informatics*, 25(2), 147–164.
2. Keys, K., Moore, D., Koga, R., Lagache, E., Tesch, M., & Claffy, K. (2001). The architecture of CoralReef: An Internet traffic monitoring software suite. In *PAM2001, Workshop on passive and active measurements*, RIPE, Citeseer.
3. Karagiannis, T., Broido, A., Brownlee, N., Claffy, K. C., & Faloutsos, M. (2004). Is P2P dying or just hiding? In *Global telecommunications conference, 2004. GLOBECOM'04. IEEE* (Vol. 3, pp. 1532–1538). IEEE.
4. Sen, S., Spatscheck, O., & Wang, D. (2004). Accurate, scalable in-network identification of P2P traffic using application signatures. In *Proceedings of the 13th international conference on World Wide Web* (pp. 512–521). ACM.
5. Dusi, M., Gringoli, F., & Salgarelli, L. (2011). Quantifying the accuracy of the ground truth associated with Internet traffic traces. *Computer Networks*, 55(5), 1158–1167.
6. Karagiannis, T., Papagiannaki, K., & Faloutsos, M. (2005). Blinc: Multilevel traffic classification in the dark. In *ACM SIGCOMM computer communication review* (Vol. 35, pp. 229–240). ACM.
7. Kim, H., Claffy, K. C., Fomenkov, M., Barman, D., Faloutsos, M., & Lee, K. (2008). Internet traffic classification demystified: Myths, caveats, and the best practices. In *Proceedings of the 2008 ACM CoNEXT conference* (p. 11). ACM.
8. Dainotti, A., Pescapé, A., & Claffy, K. C. (2012). Issues and future directions in traffic classification. *IEEE Network*, 26(1), 35–40.
9. Bermolen, P., Mellia, M., Meo, M., Rossi, D., & Valenti, S. (2011). Abacus: Accurate behavioral classification of P2P-TV traffic. *Computer Networks*, 55(6), 1394–1411.
10. Foremski, P., Callegari, C., & Pagano, M. (2014). DNS-Class: Immediate classification of IP flows using DNS. *International Journal of Network Management*, 24(4), 272–288.
11. Finamore, A., Mellia, M., Meo, M., & Rossi, D. (2010). KISS: Stochastic packet inspection classifier for udp traffic. *IEEE/ACM Transactions on Networking*, 18(5), 1505–1515.
12. Foremski, P., Callegari, C., & Pagano, M. (2014). Waterfall: Rapid identification of IP flows using cascade classification. In *Computer networks* (pp. 14–23). Springer.
13. Kuncheva, L. I. (2004). *Combining pattern classifiers: Methods and algorithms*. New York: Wiley.
14. Foremski, P., Callegari, C., & Pagano, M. (2015). Waterfall traffic identification: Optimizing classification cascades. In *Computer networks* (pp. 1–10). Springer.
15. Fiadino, P., Bär, A., & Casas, P. (2013). HTTPTag: A flexible on-line HTTP classification system for operational 3G networks. In *International conference on computer communications, 2013. INFOCOM'13*. IEEE.

16. Adami, D., Callegari, C., Giordano, S., Pagano, M., & Pepe, T. (2012). Skype-Hunter: A real-time system for the detection and classification of Skype traffic. *International Journal of Communication Systems*, 25(3), 386–403.
17. Korczynski, M., & Duda, A. (2014). Markov chain fingerprinting to classify encrypted traffic. In *INFOCOM, 2014 Proceedings IEEE*. IEEE.
18. Duda, R. O., Hart, P. E., & Stork, D. G. (2012). *Pattern classification*. New York: Wiley.
19. Dainotti, A., Pescapé, A., & Sansone, C. (2011). Early classification of network traffic through multi-classification. In *Traffic monitoring and analysis* (pp. 122–135).
20. Alpaydin, E., & Kaynak, C. (1998). Cascading classifiers. *Kybernetika*, 34(4), 369–374.
21. Chellapilla, K., Shilman, M., & Simard, P. (2006). Optimally combining a cascade of classifiers. *Proceedings of SPIE*, 6067, 207–214.
22. Abdelazeem, S. (2008). A greedy approach for building classification cascades. In *Seventh international conference on machine learning and applications, 2008. ICMLA'08* (pp. 115–120). IEEE.
23. Gringoli, F., Salgarelli, L., Dusi, M., Cascarano, N., Risso, F., & Claffy, K. (2009). Gt: Picking up the truth from the ground for internet traffic. *ACM SIGCOMM Computer Communication Review*, 39(5), 13–18.
24. Bujlow, T., Carela-Español, V., & Barlet-Ros, P. (2015). Independent comparison of popular DPI tools for traffic classification. *Computer Networks*, 76, 75–89.
25. Carela-Español, V., Bujlow, T., & Barlet-Ros, P. (2014). Is our ground-truth for traffic classification reliable? In *Passive and active measurement* (pp. 98–108). Springer.
26. Carela-Español, V., Barlet-Ros, P., Cabellos-Aparicio, A., & Solé-Pareta, J. (2011). Analysis of the impact of sampling on NetFlow traffic classification. *Computer Networks*, 55(5), 1083–1099.
27. Alcock, S., & Nelson, R. (2012). Libprotoident: Traffic classification using lightweight packet inspection. WAND Network Research Group, Technical Report.



Paweł Foremski is a Ph.D. student at the Institute of Theoretical and Applied Informatics of the Polish Academy of Sciences (IITiS PAN). He received his M.Sc. Eng. degree in Informatics (Macrocourse) in 2011 from the Silesian University of Technology in Gliwice, Poland. He works as a research assistant at IITiS PAN and as an IT consultant. His professional interests include Computer Networks, IPv6, DNS, Traffic Classification, Open Source, and Machine Learning.



Christian Callegari received the B.E. and the M.E. degrees in telecommunications engineering and the Ph.D. degree in information engineering from the University of Pisa, Italy, in 2002, 2004, and 2008, respectively. Since 2005, he has been with the Department of Information Engineering at the University of Pisa. In 2006/2007, he was a visiting student research collaborator at the Department of Computer Science at ENST Bretagne, France, and in 2014 he was a visiting researcher at Eurecom in France. Dr. Callegari is currently a researcher at RaSS National Laboratory (CNIT) and teaching assistant at the University of Pisa. Moreover he has given lectures in the framework of several Ph.D. courses (both at national and international level) and he has also given several tutorials about network security in leading international conferences. His research interests are mainly in the area of network security, with focus on Anomaly Detection and distributed architecture for security monitoring and privacy aware data exporting and processing. Moreover, he has co-authored more than 80

papers presented in leading international journals and conferences.



Michele Pagano received laurea (cum laude) in Electronics in 1994 and a Ph.D. in Information Engineering in 1998, both at University of Pisa. From 1997 to 2007 he has been Researcher at the Department of Information Engineering of the same university, and then became associate professor (confirmed in 2010). Currently he is the official instructor of the courses of Telematics, Performance of Multimedia Networks, Network Security and Architectures, Components and Network Services. In 2006, in collaboration with Prof. Vaton, he gave a Ph.D. Course on “IP traffic characterization, data analysis and statistical methods: Bayesian Methods in Teletraffic Theory”. Furthermore, he gave lectures on Network Performance Analysis in different Polish and Russian universities. His research interests are related to statistical traffic characterization and network performance analysis, statistical traffic classification, anomaly detection, security issues in distributed architectures and Green Networking. He has co-authored around 200 papers published in international journals and conference

proceedings. He has been involved in the activities of the NoE Euro-NGI (Design and dimensioning of the Next Generation Internet) and in several national and international projects, being the local coordinator for the 2006 PRIN RECIPE (Robust and Efficient traffic Classification in IP nEtworks) and the 2008 PRIN EFFICIENT (Energy eFFicient teChnologies for the Networks of Tomorrow). In 2006/2007 he has been supervisor of Dr. Marchenko in an INTAS grant and in 2009/2012 he has been the principal investigator in two inter-university cooperation projects with PetrSU, PFUR and TvSU. Finally, in 2011–2013 he has been the supervisor of Pawel Foremski in the framework of the project “Multilevel traffic classification in the Internet”, funded by the Polish National Science Centre.