# Design of DIFFUSE v0.4 – DIstributed Firewall and Flow-shaper Using Statistical Evidence

Sebastian Zander, Grenville Armitage Centre for Advanced Internet Architectures, Technical Report 110704A Swinburne University of Technology Melbourne, Australia szander@swin.edu.au, garmitage@swin.edu.au

Abstract-In recent years a growing number of researchers investigated the performance of machine learning based traffic classification using statistical properties - classification techniques that do not require packet payload inspection. Such techniques assist Internet Service Providers to work within any legal or technical limitations on direct payload inspection. Potential new applications include automated 'market research', automated traffic prioritisation, and Lawful Interception. For many of these new applications a de-coupling between the flow classification and subsequent flow treatment, such as blocking or shaping, is highly desirable. We developed DIFFUSE - an extensions for an existing packet filter that provide MLbased traffic classification based on statistical properties and de-couple flow classification from flow treatment. This report describes the selection of the existing packet filter extended, the design of the overall architecture and key components, as well as the machine learning techniques supported. This report is an updated version of tech report 101223A [1].

Index Terms—Statistical Flow Classification, Machine Learning, Quality of Service, Traffic Prioritisation

#### I. INTRODUCTION

During recent years a body of research emerged around the identification and classification of traffic flows based on statistical properties (features) and in particular the application of Machine Learning (ML) techniques to generate such classifiers [2]. Statistical properties, such as distributions of packet sizes or inter-packet arrival times, can be calculated without accessing packet payloads (payload inspection). Such techniques assist Internet Service Providers (ISPs) to work within any legal or technical limitations on direct payload inspection. Potential new applications include automated 'market research', characterising traffic for Lawful Interception [3], or automated prioritisation of real-time traffic [4].

For many of these new applications a decoupling between flow classification and subsequent flow treatment (the actions performed on flows), such as blocking or shaping, is highly desirable. For example, a single high performance classifier near the core of an ISP network may control multiple low-power nodes near the network edge (perhaps embedded within Asynchronous Digital Subscriber Line or Cable modem gateways) so that centralised traffic classification can automatically modify the Quality of Service (QoS) treatment experienced by packets at the network edge. This decoupling also enables potentially computationally intensive per-flow statistics calculations to be offloaded from the packet forwarding path.

Common open-source packet filters that combine firewall and traffic shaping (such as IPFW [5], PF [6], Netfilter [7] and others) currently do not use traffic statistics, instead relying on direct inspection of packets passing through the filtering node's local interfaces. Furthermore, these filters couple the flow classification and treatment tightly, i.e. the actions are executed locally immediately after the flow classification.

In the DIFFUSE project [8] we are designing and developing extensions for an existing packet filter that provide ML-based traffic classification based on statistical properties and decouple flow classification and treatment. In our architecture there are *classifier nodes* that classify traffic flows and then instruct *action nodes* via a *control protocol* to carry out actions for the classified flows.

In this report we describe the design of the system and its key components. To avoid 'reinventing the wheel' our system will be based on an existing packet filter. As main development platform we selected the FreeBSD operating system since it is often used for building firewalls and/or traffic shapers, and a number of existing packet filters run on FreeBSD. This report is an updated version of a tech report describing DIFFUSE version 0.1 [1] and obsoletes the older report. The report is organised as follows. First we define fundamental terms and concepts in Section II. In Section III we explain which FreeBSD packet filter we chose as platform to develop and demonstrate DIFFUSE. In Section IV we describe the overall architecture of the system and its key components, and show example scenarios illustrating how the system could be used.

In Section V we describe the extended command language that enables the configuration of classifier and action nodes. In Section VI we describe the software design of classifier and action nodes. In Section VII we describe the design of the control protocol used to exchange data between classifier and action nodes.

In Section VIII we outline how classifier models can be created, and how DIFFUSE can be used for offline experiments. In Section IX we describe the initial choice of ML techniques supported. Since our system is designed to be flexible other ML techniques can be added in the future. Section X concludes the report.

## II. DEFINITIONS

First we define some fundamental terms and concepts used throughout the report.

## A. Flows

A *flow* is a number of consecutive packets that have the same values for a defined set of packet header fields within a certain time frame. The set of packet header fields is usually the commonly used 5-tuple (source and destination IP addresses, source and destination ports, protocol), but it could be a different set of fields (such as only the source and destination IP addresses).

For connection-oriented protocols (like TCP) the flow start and end is usually marked by the establishment and teardown of a connection. For non connection-oriented protocols (like UDP) the first packet seen marks the start and no packets arriving for a certain duration (flow timeout) marks the end of the flow.

A *unidirectional* flow is a flow where packets flow only in one direction (e.g. only packets matching a 5tuple), whereas a *bidirectional* flow is packets flowing in both directions (e.g. all packet matching a 5-tuple and the same 5-tuple with source/destination addresses and ports reversed).

A *sub-flow* is part of a flow. For our purposes a sub-flow is a sliding window of n consecutive packets within a unidirectional flow (as in [9]).

Bidirectional flows have two directions which we refer to as *forward* and *backward*. For connection-oriented protocols (and if the initial handshake can be observed) packets from the originator of the connection are going in the forward direction, and packets from the other end are going in the backward direction. For connection-less protocols or when the handshake could not be observed the first packet defines the forward direction.

If a rule defines a source or a destination (by specifying a match pattern, e.g. matching against the source IP address), packets matching the pattern are considered to flow forward, whereas packets that are going in the reverse direction are considered to flow backward.

# B. Features

Previous work usually used the two level hierarchy of *features* and *feature sets*, where a feature is a characteristic of a flow or sub-flow (such as the mean packet length) and a feature set is a number of different features. For the DIFFUSE architecture we extended this hierarchy. *Feature statistics* are statistics of a series of feature values (such as the minimum, mean or maximum), features are characteristics of flows, sub-flows or packets (such as packet length or inter-arrival times) and feature sets are a number of features (as before).

The main reason for this three-level hierarchy is that DIFFUSE v0.4 supports different independent feature modules, but for performance reasons different statistics of the same feature are part of the same module.

# III. CHOICE OF FIREWALL

Here we discuss our choice of the existing packet filter we extended. Since DIFFUSE v0.4 is based on FreeBSD we have a choice between three packet filters: IP Firewall (IPFW) [5], IPFilter (IPF) [10], [11] and Packet Filter (PF) [6]. We compare these based on a number of criteria.

#### A. Functionality

All three packet filters support the basic functions of filtering based on network and transport layer information, network address translation, logging etc. IPFW and PF can tag packets for implementing policy-based rules and have interfaces to traffic shapers that can queue and prioritise packets. IPFW mostly uses Dummynet [12] but also has an interface to ALTQ [13], whereas PF uses ALTQ. IPFW/Dummynet can also be used to emulate certain network link conditions by limiting link capacity, emulating delay and packet loss etc.

PF has more advanced functionality than IPFW and IPF, most notably the ability to implement redundant firewalls (state transfer and fail-over protocol [6]), load-balancing, logging to tcpdump files, and filtering on

operating system fingerprints. However, IPFW now also supports tables and in-kernel NAT and has reduced the functionality gap to PF.

DIFFUSE needs packet queuing and prioritising support, which rules out IPF. While the advanced functions of PF are nice they are not really required for DIFFUSE.

# B. Portability

IPFW has been developed and used in FreeBSD over many years, is the network firewall in MacOSX, and has recently been ported to Linux and Windows [12]. IPF runs on the BSD family (FreeBSD, OpenBSD, NetBSD) as well as on Solaris, HP-UX, IRIX and Linux. PF is the main firewall of OpenBSD, and it has been ported to FreeBSD and NetBSD.

As far as portability is concerned IPFW and IPF are the best options for DIFFUSE. PF falls behind as the number of operating systems it runs on is very limited.

# C. Support

IPFW is the FreeBSD sponsored firewall; it is authored and maintained by FreeBSD volunteer staff members. IPF was the main packet filter of OpenBSD, before it was replaced by PF in 2001. Given that we have chosen FreeBSD as development platform in terms of future support IPFW is most promising. The PF firewall comes second being OpenBSD's official firewall and given that it is part of the FreeBSD source tree. IPF is also part of the FreeBSD source tree.

However, the PF sources part of FreeBSD are always lagging behind the latest OpenBSD version. For example, the PF version in FreeBSD-9.0-current checked out in July 2010 are the PF sources from March 2007. Similarly, the IPF sources in FreeBSD lack one major version behind the latest release. For example, the IPF version in FreeBSD-9.0-current checked out in July 2010 are the IPF sources from October 2007 (version 4.1) and much older than the latest release from May 2010 (version 5.1). For IPF this is less of a problem because the IPF sources are released independently of FreeBSD and should compile on all supported operating systems.

All three packet filters are actively maintained and used. For DIFFUSE IPFW and PF have a slight edge as they are the main packet filters of FreeBSD/OpenBSD.

# D. Performance

Measuring the performance of packet filters is not straightforward, as the performance depends heavily on the scenario, i.e. the actual firewall rules and network traffic. Nevertheless, one can compare different packet filters in particular scenarios. Previous work compared IPF, PF and Linux Netfilter [14], [15]. According to these studies PF performs similar to IPF and whether PF and IPF perform better or worse than Netfilter depends on the scenario. We were unable to find a study of IPFW or a recent performance comparisons of all three firewalls.

The performance of the packet filter should be reasonably good, but this criteria is not of very high importance for DIFFUSE. Since IPFW, IPF, and PF are are all deployed we believe they all provide sufficient performance in practice.

# E. Usability

The rule set language of PF is better designed than the languages of IPF and IPFW, which both look somewhat organically grown. The structure of PF and IPF rulesets differs from IPFW (and Linux Netfilter). By default for IPF and PF the last rule that matches determines the action, whereas for IPFW (and Netfilter) the first rule that matches determines the action. This makes it more difficult to convert IPF and PF rulesets to IPFW rulesets and vice versa.

While the rule language of PF is better designed, the languages of IPF and IPFW are still logical, easy to read and use. As far as usability is concerned all three firewalls are adequate for DIFFUSE.

# F. Extensibility

IPFW, IPF and PF have nicely written user documentation, but for all three there is not much developer documentation. None of them has a fully modular framework that can be extended easily. However, IPFW's Dummynet now has a modular framework for adding packet schedulers. All three packet filters have nicely written code, but IPFW stands out because it also has a lot of useful comments inside the code, whereas comments in IPF and PF code are rather sparse. Furthermore, here at CAIA we have some in-house expertise for extending IPFW. Hence, IPFW wins this category.

# G. Decision

IPF does not have functions for packet queuing and prioritisation and hence cannot be used. We chose IPFW over PF because IPFW supports the three arguably most popular operating systems (FreeBSD, Linux, Windows) and it appeared to be easier to extend than the others due to well documented code, relatively modular structure and existing in-house expertise.

#### **IV. SYSTEM DESCRIPTION**

# A. Architecture

In order to provide ML-based traffic classification and decouple the classification from the subsequent action the DIFFUSE system has several key components:

- A *Classifier Node* (CN) computes statistical features from flows identified by their 5-tuple and classifies them based on local machine-learning rules.
- An *Action Node* (AN) performs configured actions (block, redirect, rate shape, etc.) on packets belonging to flows that have been classified by a local or remote CN.
- An IP-layer *Control Protocol* (CP) between CNs and ANs to enable real-time coordination, such as alerting ANs when to start and stop acting on identified flows.
- An extended set of *Packet Filter Rules* (PFRs) to express ML-based traffic matching based on statistical attributes at CNs and specify the actions to be taken by nominated ANs.

A CN records flow identification information (5-tuple) and computes flow characteristics, such as packet length and inter-arrival time statistics. It continuously compares the statistics of observed flows to a configured set of rules and uses this information to generate traditional header-only inspection rules for ANs. When a flow (flow X) matches a statistical rule, the CN passes the flow's 5-tuple and class to AN(s) to actually instantiate the flow class' associated action. The action is then applied to all subsequent packets belonging to flow X. The rule is removed from the AN(s) once flow X has stopped.

CNs and ANs automatically establish IP based control links via a CP to share information as matching flows come and go. CNs and ANs are different logical entities, but they can be co-located on the same physical network device. For example, a traditional packet filter combines them in a single device. In this case the control link is inside the host.

Small networks with a few CNs and ANs can be configured manually by creating and distributing rulesets. In large networks comprising many CNs and ANs it is desirable to have a management system that automatically translates network policies into rulesets that are distributed to CNs and ANs. Such a management system is out of scope of this document, but it can be designed and build on top of our developed system in future work.

#### B. Classifier Node

A CN consists of an extended packet Filter/Classifier in kernel space and a userspace daemon process (called *Exporter*) that exports the flow specification (5-tuple), class and (optionally) an action to the AN(s) via the CP (see Figure 1). We call this information *flow rules*.

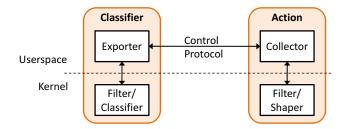


Figure 1. Classifier node and action node components, and control protocol

The extended packet filter computes statistical features for packet flows. These features can be used directly as patterns for matching packets or as input for an MLbased classifier that assigns classes to packets based on these features.

The feature computation and classification is done on a per-packet basis inside the kernel to maximise performance. The sending of flow rules to remote ANs is done by a userspace daemon, because this task is less performance critical<sup>1</sup> and a userspace daemon is easier to develop, test, and port to other operating systems (OSs), and it has access to a much larger set of functionality via libraries.

The Exporter needs access to flow information generated by the classifier. Since the IPFW control interface (based on socket options) does not allow unsolicited messages from kernel to userspace and frequent polling of the in-kernel classifier is impractical, a separate interface (UDP socket) is used to send the flow information from classifier to Exporter.

Section VI describes the CN design in more detail.

# C. Action Node

The AN consists of an userspace daemon (called *Collector*) that listens for flow information from CNs and configures the packet filter and traffic shaper accordingly using the existing IPFW configuration interface(s). The Collector consists of a frontend and a backend. The frontend handles the control protocol communication

<sup>&</sup>lt;sup>1</sup>We assume the number of simultaneous active flows that trigger remote actions is typically only a few thousand and there is a limit on how often actions are updated.

and manages the addition/removal of flow rules stored in an internal database. The backend is responsible for generating IPFW traffic filter/shaper rules based on the flow rules in the database.

The advantage of a userspace daemon is that no kernel code needs to be modified, again easing development, testing, porting to other OS and access to userspace libraries. Furthermore, in the future different firewalls or traffic shapers can be supported easily by modifying or extending the backend of the Collector. Besides packet filter or traffic shaper specific actions, other actions could be implemented, such as logging of classified traffic in a database.

Section VI describes the AN design in more detail.

#### D. Ruleset operations

On the CN the rule language of IPFW has to be extended to allow the specification of features to be computed, use of feature values in match patterns, use of ML classifiers, and the configuration of remote ANs (see Section V). On the AN only the existing packet filter and traffic shaping functionality is used and no rule set language modifications are required.

However, in addition to the extended rule set language we need new commands to configure the Exporter and Collector (see Section V).

#### E. Control protocol operations

CNs send "*add rule*" messages (ARMs) to ANs (see Figure 3), which contain (partial) rules that have a match part (flow specification), one or more class tags, and optionally an action part. Our notion of ARMs includes the updating of existing rules. If a CN sends an ARM to an AN with the same flow specification but different action as a rule sent previously, the AN must replace the previous rule with the new rule.

If a CN is configured such that it does not send actions in ARMs, the receiving AN(s) must be configured so that they have a list of flow classes and associated actions. In this case the AN(s) will determine the actions based on the classes identified by the CN. If the AN(s) are configured with such action lists, the configured actions always overrule any actions specified in ARMs.

Rules can have a *timeout* which will cause ANs to remove rules after a specified duration has elapsed since the rule became active (*rule timeout*), or when no packets have matched the rule for the specified duration (*flow timeout*). This is shown in Figure 2. If rule timeouts are used CNs need to periodically refresh rules (so that long running flows are properly handled).

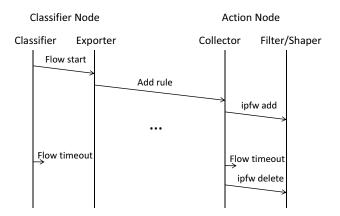


Figure 2. Rule creation and timeout based rule removal

CNs can also send explicit "*remove rule*" messages (RRMs) to ANs (see Figure 3). An AN will remove all rules that match the rule specification in the message. If the rule specification is broad, e.g. only the protocol is specified, this may trigger the removal of many rules. If one wants to remove exactly one rule, the flow specification in the RRM must be specified exactly as in the ARM.

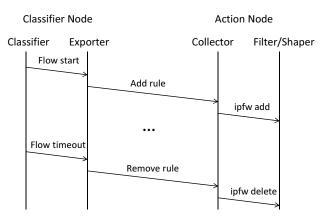


Figure 3. Rule creation and explicit rule removal messages

The use of flow timeouts has advantages over explicit RRMs. Firstly, they save network capacity as the number of messages can be reduced. Secondly, they can also prevent control loops that may occur otherwise because actions like blocking or shaping affect packet flows and their characteristics and hence the decisions of CNs.

Imagine the following scenario. A CN identifies a flow to block, triggers an AN, and the AN blocks the flow. The flow times out at the CN triggering a RRM from the AN, which unblocks the flow. The CN identifies the flow again and so on. If the AN handles the flow timeout locally (using a timeout provided by the CN), the rule will be active as long as packets are matching and then it will timeout. This approach presupposes that the flow times out at the AN before it times out at the CN, but the CN can ensure this by controlling the AN's timeout value. This approach also covers the case when the rule is (prematurely) removed due to resource constraints at the AN.

However, flow timeouts may not be available everywhere (they may not be efficient to implement on all devices) and then RRMs must be used. But the delivery of RRMs cannot be guaranteed in all cases, e.g. CNs may crash. To avoid rules living forever in the AN, the AN purges old rules based on last recently used information (flow timeout). Should the AN run out of memory it is up to the AN to decide what to do, i.e. what rules to purge, unless it was advised by the CN what to do. For example, if the CN attached priorities to rules, then the AN must purge rules according to their priority.

In our prototype implementation by default the CN uses explicit RRMs, and the AN uses flow timeouts to time out rules. Explicit RRMs can be turned off, then both the CN and AN use flow timeouts.

The CN classifies flows all the time (in the extreme case for each packet arriving), but the CN will only notify AN(s) for new or changed flows (same flow specification but different class or action). Note, that a new flow may have the same flow specification as a previous flow if the previous flow has ended before the new flow started. Furthermore, there are options to further limit the sending of rules, see Section V-H.

Section VII describes the details of the protocol.

#### F. Example Scenarios

We illustrate the DIFFUSE v0.4 architecture in an example scenario, where the ISP differentiates a customer's traffic into real-time and non-real-time traffic and subsequently uses this information to prioritise the real-time traffic. Figure 4 shows the customer and the ISP network. A CN with a rule database is located on or connected to an edge router inside the ISP's network. Two ANs are located on the ISP's edge router and customer's router.

During operation the system does the following. The CN continuously classifies traffic flowing between the customer and ISP networks based on statistical characteristics and stored rules. For each new real-time flow detected, the CN sends the flow's 5-tuple, class and action to the ANs. The ANs then create a new rule for the real-time flow that will prioritise its traffic over non-real-time flows. After the real-time flow has stopped, the rule is removed from the ANs.

#### V. DESIGN OF COMMAND SET EXTENSIONS

Here we describe the DIFFUSE v0.4 extended IPFW command set used to configure CN and AN.

#### A. Notation

We use the following notation based on ABNF [16]. Bold typewriter font identifies parameter names (terminal symbols) and *italic* typewriter font identifies parameter values chosen by the user (non-terminal symbols that do not contain spaces). Symbols in double quotes "" are also terminal symbols.

Normal typewriter font identifies nonterminal symbols that are broken down into parameter names and values at a later point. Parameters and values in square brackets [] are optional, a slash / defines alternatives, and round brackets () are used for grouping.

A preceding n\*m means a symbol or group is repeated a minimum n and a maximum m times. If n is zero it is omitted, and if m is infinity it is omitted as well. This allows short forms such as \* for 0-infinity or 1\* for 1-infinity.

#### B. Existing IPFW rules

First we define existing IPFW rules and a number of symbols for parts of existing IPFW rules that we later use in the definitions of extended rules (see Figure 5). The symbol ipfw-rule-id is the optional rule number, the rule set number and it also includes IPFW's match probability. The symbol ipfw-log-altq-tag comprises the log, ALTQ and tag options. ipfw-action is one of the actions executed when the pattern part of the rule matches (note that ... is a placeholder for the other actions not shown here [5]). The symbol ipfw-patterns describes the patterns that are used to match a packet and options describes all possible options, such as keep-state, tagged [5].

#### C. Configure, delete and show features

Figure 6 shows the command to configure a feature. The argument feature-name is a user-defined name for the feature. The argument module-name is the name of the feature as defined in the feature's implementation (basically an implemented feature module is the class and a configured feature is an instance of the class). The feature's kernel module must be available, i.e. loaded into the kernel.

The symbol options stands for further options provided by the feature module. Each option is either a flag (enabling or disabling a property of a feature) or a parameter name followed by an argument. Our prototype

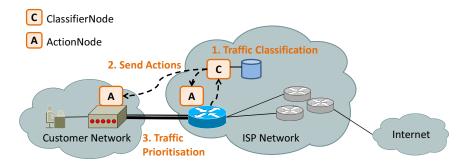


Figure 4. Using the DIFFUSE architecture for distributed traffic prioritisation

```
ipfw-rule = ipfw-rule-id ipfw-action ipfw-log-altq-tag ipfw-patterns
ipfw-rule-id = [rule-number] [set set-number] [prob probability]
ipfw-log-altq-tag = [log [logamount number]] [altq queue] [(tag / untag) number]
ipfw-action = (allow / deny / count / ...)
ipfw-patterns = proto from source to destination [ipfw-options]
ipfw-options = (keep-state / tagged tag-list / ...)
```

Figure 5. Existing IPFW rules

```
ipfw feature feature-name config module module-name [options]
options = *((option-flag / (option-name option-value))))
```

```
ipfw feature ( delete / show ) feature-name
```

Figure 6. Configure, delete and show features

implementation creates default instances for all existing features with feature-name set to module-name. When a feature has no options or the default options are deemed acceptable the default feature instances can be used straight-away.

A feature can be viewed or deleted with the commands shown in Figure 6. The reserved feature name all can be used to show or delete all defined features.

## D. Current features and their options

There are two types of features: unidirectional and bidirectional. Unidirectional features compute statistics for unidirectional flows. For bidirectional flows these are computed separately for each direction. Bidirectional features compute statistics over both directions of bidirectional flows (or only one direction of unidirectional flows).

The initial prototype has the features shown in Table I. The "plen" feature provides minimum, mean, maximum, standard deviation and sum of packet-length statistics. The "iat" module provides minimum, mean, maximum, and standard deviation of inter-arrival time statistics. As the names suggest, the "plenbd" and "iatbd" features are the bidirectional versions of "plen" and "iat" and provide the same statistics over bidirectional flows. The "pcnt" feature provides packet count statistics. The "skype" feature provides mean packet length, packet ratio and absolute two-packet length difference statistics tailored to detect Skype traffic.

All of our initially implemented features provide the window and partial-window options. The option window defines the window in packets over which a feature is computed. By default features are only computed for full windows, unless partial-window is set. Any match for a feature statistic that has not been computed yet due to a partial window will fail. Note that a feature does not necessarily need to be computed over windows of packets, i.e. there could be features computed for single packets only.

By default the packet length computed by "plen" is the packet length of an IP packet including the IP header. If the option ipdata-len is used, the packet length computed is the length of an IP packet excluding the IP header. If the option payload-len is used the length computed is the payload length (UDP or TCP data). The "iat" (inter-packet arrival time) feature has an option

Table I				
FEATURES AND THEIR OPTIONS				

Feature	Туре	Options			
plen	uni	window size, partial-window, payload-len, ipdata-len			
plenbd	bidir	window size, partial-window, payload-len, ipdata-len			
iat	uni	<pre>window size, partial-window, accurate-time</pre>			
iatbd	bidir	window size, partial-window, accurate-time			
pcnt	bidir	window size, partial-window			
skype	bidir	window size, partial-window, payload-len, ipdata-len			

accurate-time. If set the computed timestamps are more accurate, at the cost of more processing time.

 Table II

 CLASSIFIER ALGORITHMS SUPPORTED BY DIFFUSE v0.4

Figure 7 shows examples how features are defined.

#### E. Compute and use features for matching

Features are computed and used for matching as shown in Figure 8. Such rules compute the specified features for all sub-flows that match <code>ipfw\_patterns</code>. Once the feature statistics are available they can be used for matching and when all IPFW patterns and feature patterns match the rule's action will be executed.

A feature-list is a comma-separated list of one or more feature names, specifying the features to be computed. Note that if one or more feature-test are present, all features used in the tests will be implicitly added to feature-list if not already present. This means that normally one does not need to specify feature-list.

By default (without unidirectional) unidirectional features are computed separately for both directions of bidirectional sub-flows and feature tests can be used with an optional direction attribute. If unidirectional is specified, features are computed for unidirectional sub-flows and hence a rule matches unidirectional sub-flows ("fwd" or "bck" are not used in feature matches). Bidirectional features are computed over both directions of bidirectional flows and hence "fwd" or "bck" are also not used in feature matches.

If feature matches are used but features is not specified an implicit features is generated, listing all features used in the tests and producing bidirectional flows. Note that if unidirectional is not set rules match against features computed in both directions and if a rules matches the action will be executed for packets in *both* directions as well. With unidirectional a rule's action is only applied to unidirectional flows.

The ipfw commands delete, flush, list, show, zero and resetlog work with feature-enabled rules as before. Figure 9 shows example commands for using feature matches.

Classifier algorithm-name	Description
c4.5	C4.5 decision tree classifier
nbayes	Naïve Bayes classifier

#### F. Use machine learning classifier

An ML classifier takes a set of features of a sub-flow F(n = |F|) as input and returns a classification c, where c is exactly one class out of the set of classes C(m = |C|). A classifier is configured as shown in Figure 10.

The argument classifier-name is the name of the classifier. The argument algorithm-name defines the classification algorithm to use and must be a name of a classifier module present (loaded into the kernel). Table II lists the module names of the currently implemented algorithms (C4.5 and Naïve Bayes), a more detailed description of the algorithms is in Section IX.

The argument of model defines the classifier model to use for the classification. The model file is a text file created by a modified version of the Waikato Environment for Knowledge Analysis (WEKA) [17] (see Section IX-C). The model file is parsed and the resulting classifier is setup in the kernel.

The model file contains the list of statistics needed, but the names may not be valid DIFFUSE names. Hence the parameter use-feature-stats allows to specify the list of features statistics the classifier uses for classifying flows. This parameter is also handy if one wants to use the same classifier model with differently generated feature statistics, for example the same statistics but generated over different window sizes. If used this parameter always overrules the statistics listed in the model file.

The argument of use-feature-stats must list the same number of feature statistics in exactly the same order they were used when building the classifier. Feature statistics are listed in the same format as for feature matches explained above. As for feature matches, ipfw feature myplen config plen window 10
ipfw feature myiat config iat window 20 accurate-time

Figure 7. Feature configuration examples

ipfw add ipfw-rule-id ipfw-action ipfw-log-altq-tag ipfw-patterns [features feature-list]
[1\*feature-test] [unidirectional] [every / once / sample packets]

feature-list = feature-name[\*(","feature\_name)]
feature-test = [(fwd / bck)"."]feature-statistic"."feature-name(<=/</=/>)value

Figure 8. Matching with features

ipfw feature myplen config plen window 25 ipfw add deny udp from any to any features myplen min.myplen<100 bidirectional ipfw add deny udp from any to any fwd.min.myplen<100 # same as above rule</pre>

Figure 9. Feature matching examples

```
ipfw mlclass classifier-name config algorithm algorithm-name model file-name
[use-feature-stats feature-stat(n-1)*(","feature-stat)] [class-names name(m-1)*(","name)]
[confirm number]
feature-stat = [(fwd / bck)"."]feature-statistic"."feature-name
```

ipfw mlclass ( delete / show ) feature-name

Figure 10. Configure, delete show classifier

for unidirectional features and bidirectional flows the direction is assumed to be forward (fwd) if not specified.

The parameter class-names can be used to overwrite the class names specified in the model file. The class names are used to match packets as shown in Figure 11. One can either use rules with the new mlclass action and use IPFW tags, or use rules with the new match-if-class option, or a combination of both.

For rules with mlclass action or match-ifclass option(s) feature lists are generated implicitly in the same way as for feature matches (see previous sub section). For bidirectional flows actions of matching rules will be applied to both directions of flows.

The parameter classifier\_name references a classifier configured previously. The parameter class-tags defines a list of IPFW tags, where each tag is associated to one class in C. As the result of the classification each packet is tagged with the tag configured for the class it matches. The tags can be used in subsequent rules for matching (with the tagged option of IPFW).

The new match-if-class parameter specifies a classifier name and a list of class names or indices that make a rule match. A preceding hash symbol

"#" differentiates between class names and class numbers (see Figure 11). If the current class of a packet or sub-flow matches any of the classes listed the match\_if\_class option matches.

The confirm parameter (in Figure 10) is a simple filter for the classification process. It specifies how many times a class needs to be "confirmed" before a classifier will match. For example, if confirm is set to two a match-if-class will only match if at least three consecutive packets have been classified as the *same* class. (By default confirm is set to zero.)

Mutually-exclusive options exist for controlling how often flows are (re)classified. By default a flow is (re)classified at every packet. If once is specified a flow is classified only once (when all feature statistics have been computed for the first time). With sample a flow is classified periodically at every packets packets, and with rnd-sample a flow is classified at packets selected with a uniform random probability prob (using the kernel's random() function). The previous class (if any) is assigned to packets not (re)classified. To minimise classification latency when sample or rnd-sample are used, the first packet where all fea-

CAIA Technical Report 110704A

```
ipfw add ipfw-rule-id ( mlclass classifier-name / ... ) ipfw-log-altq-tag ipfw-patterns
[class-tags tag(m-1)*(","tag)] [match-if-class class-name:class[*(m-1)(","class)]] [once /
sample packets / rnd-sample prob]
```

```
class = ( name / "#"number )
```

Figure 11. Classes-based matching

ture statistics have been computed for the first time is always classified.

A classifier can be viewed or deleted with the commands shown in Figure 10. The reserved classifier name all can be used to show or delete all defined classifiers. Note that classifier models are only shown for single classifiers, but not when all is used.

Without IPFW tags one can write rules that use an ML-based classifier as shown in Figure 12. If IPFW tags are used one can write rules as shown in Figure 13.

Flows can be classified by multiple classifiers. Hence multiple match-if-class with different class-name can be used in one rule. The standard IPFW options can be used to select what type of flows are classified. For example, if one wants to classify only UDP flows the "ip" in the example rules can be replaced by a "udp".

Note that classification is only performed once all feature statistics needed are available. For example, if the statistics have not been computed yet because windows are only filled partially and the feature option partial-window (described in Section V-D) has not been used, the flow will not be classified.<sup>2</sup> (In future work some of the classifiers could be extended to handle missing features.)

## G. Flow table

State of all flows for which features are computed is stored in a flow table. The flow table can be displayed with the show command (see Figure 14). This command shows information for all flows, such as the rule that generated the flow, packet and byte counters, the 5-tuple, a list of all computed features and their current values and a list of all class labels.

By default only active flows are shown, but if the expired parameter is specified expired flows are also shown. (Expired flows are no longer active but their state is still in the table. State is not immediately deleted when

flows expire, but only when more space is needed for adding new flows.)

All entries in the flow table can be removed (flushed) or the packet and byte counters can be zeroed by using the flush or zero commands (see Figure 14).

#### H. Export flow rules to ANs

On the CN we need to configure rules that decide what information the classifier sends to the Exporter or to remote AN(s).

Firstly, an export target needs to be configured as shown in Figure 15. The argument export-name is the name of the new export target instance. The argument of target is the destination of the flow rules. The protocol must be UDP, host is the fully qualified host name (or IP address), and port is the port number the target is listening on. The arguments action-name and action-params-val are the action name and parameters that are send for matching flows. Note that action name and parameters must specify valid IPFW actions<sup>3</sup>. Note that the receiving AN(s) may overrule these with locally specified actions.

The argument of min-batch is the minimum number of flow rules exported in one batch. Similarly, the argument of max-batch is the maximum number of flow rules exported in one batch (must be equal or larger than min-batch). Note that increasing min-batch also increases the delay for delivering flow rules.

The argument of max-delay specifies a maximum delay between the generation of flow rules and their export. Note that if max-delay is set (value larger than zero) the minimum batch size is still enforced, but the maximum batch size can now be exceeded (if at the time of exporting more rules are over the maximum delay than the size of the maximum batch).

The argument of confirm specifies how many times a flow has to be consecutively classified as the same class before flow information is exported. For example, if confirm is set to 2, information is only exported

<sup>&</sup>lt;sup>2</sup>Even with partial-window configured more than one packet may be required to compute the feature statistics. For example, the "iat" feature requires at least two packets to compute inter-arrival time statistics.

<sup>&</sup>lt;sup>3</sup>Currently, these are opaque values that are not checked.

ipfw mlclass myclass config algorithm c4.5 model /etc/ipfw/realtime.model use-feature-stats fwd.min.myplen,fwd.mean.myplen,fwd.max.myplen,bck.min.myplen,bck.mean.myplen,bck.max.myplen class-names rt,nonrt ipfw add pipe 1 ip from any to any match-if-class myclass:rt ipfw add pipe 2 ip from any to any match-if-class myclass:nonrt

#### Figure 12. Matching using the new match-if-class option

ipfw mlclass myclass config algorithm c4.5 model /etc/ipfw/realtime.model use-feature-stats fwd.min.myplen,fwd.mean.myplen,fwd.max.myplen,bck.min.myplen,bck.mean.myplen,bck.max.myplen class-names rt,nonrt ipfw mlclass myclass ip from any to any class-tags 1,2 ipfw add pipe 1 ip from any to any tagged 1 ipfw add pipe 2 ip from any to any tagged 2

Figure 13. Matching using IPFW tags

ipfw flowtable show [expired]
ipfw flowtable ( zero / flush )

Figure 14. Flow table commands

```
ipfw export export -name config target udp://host:port [action action-name] [action-params
action-params-val] [min-batch number] [max-batch number] [max-delay delay] [confirm number]
[unidirectional]
ipfw add ipfw-rule-id ( export export-name / ... ) ipfw-log-altq-tag ipfw-patterns
diffuse-patterns
```

Figure 15. Configure export target and trigger export of rules

if the class was confirmed twice (three consecutive classifications resulting in the same class).<sup>4</sup>

By default ANs treat flows as bidirectional, i.e. apply actions to both directions of a flow. Setting unidirectional instructs ANs to treat flows as unidirectional, but only if they were unidirectional at the CN as well. However, based on local configuration the receiving AN(s) may still decide to treat flows differently.

Secondly, we need to define the rules that, if they match, will send flow rules to the configured AN(s) as shown in Figure 15. A rule with the new export target will export flow rules according to the configured exporter export-name for all flows that match the rule. Figure 16 shows an example where all flows classified as real-time are exported to localhost.

The classifier in the kernel can only export information via the UDP transport protocol. If UDP is sufficient the classifier can send rules to ANs directly. However, in many cases UDP will not be appropriate, for example if reliable transport is required (see Section VII). In this case the classifier needs to send the information to the userspace Exporter via UDP, which then forwards the information to the Collector via SCTP or TCP (see Section VII).

The Exporter is configured as shown in Figure 17. By default the Exporter listens for flow rules from any (kernel) classifier on the default port 3191. The -c switch can be used to specify a particular classifier host and change the default port number<sup>5</sup>. The flow information is forwarded to a number of ANs specified as list of URLs with the -a switch. The -q switch turns off any output to stdout.

Note that not only the 5-tuple describing the flow, the class tags and the action is exported to ANs. A variety of other data is sent as well, such as a bidirectional flag that specifies if actions should be executed for both

<sup>&</sup>lt;sup>4</sup>Typically, this option is only used for classifiers configured without using the classifier confirm option. However, it can always be used to more aggressively filter for flow-rule exports than for local rules.

<sup>&</sup>lt;sup>5</sup>Typically the Exporter runs on the same host as the kernel classifier, but it could run on a different host.

ipfw export myexp config target udp://localhost min-batch 1 max-batch 5 ipfw add export myexp ip from any to any match-if-class myclass:rt

Figure 16. Export flow rules example

```
ipfw_exp [-c host:port] [-a list-of-urls] [-q]
list-of-urls = url*(","url)
url = (udp / tcp / sctp)"://"host":"port
```

Figure 17. Userspace exporter configuration

directions of bidirectional flows, rule timeouts and so on (see Section VII).

#### I. Listen to remote CNs

On the AN we need to configure the Collector to listen to remote Exporter(s) as shown in Figure 18. The parameters -s, -t and -u specify on which SCTP, TCP, UDP ports the Collector is listening (at least one of these must be specified). The -n switch turns off the IPFW rule generation (useful for testing as non-root). The -q switch turns off any output to stdout. The -r switch specifies the IPFW rule number space used by rules generated by the Collector (default 1000–2000). The Collector will create as many IPFW rules as fit into this space. Note that it is the users responsibility to ensure that the range specified is available.

The -c switch specifies a file that defines a mapping between classes and actions (*actions file*). If flow rules are received with one of the classes specified in the actions file, the specified actions will always overrule any actions given by the CN. The syntax of the actions file is shown in Figure 19. Figure 20 shows an example actions file.

In principle the Collector is independent of the firewall or traffic shaper used to treat flows. However, currently the Collector can only be used with IPFW. Also note that currently prototype action names and parameters are opaque values, which are not checked by the Collector.

The Collector has its own flow table. For each rule received from a CN via an ARM the Collector checks if the rule is already present in the table. If a flow rule is present with same flow specification and action, the collector only updates the timeouts (if any). If a flow rule is present with the same flow specification but different action and there is no actions file, the Collector replaces the old rule with the new rule and updates timeouts (if any). If no rule is present with the same flow specification the collector inserts the new rule into the table. If a new rule was inserted or an existing rule was updated the Collector will create a new IPFW rule or update an existing IPFW rule by using the IPFW command line interface. Removal of rules occurs upon timeout or explicit RRM. In both cases the Collector removes the rule from its internal database and then removes the IPFW rule via IPFW's command line interface.

Figure 21 shows an example for configuring an Exporter and Collector.

VI. DESIGN OF CLASSIFIER AND ACTION NODE

Now we describe the software design of CN and AN.

#### A. Classifier Node

1) Overview: Figure 22 shows the main building blocks of the CN. Inside the kernel there is a new DIFFUSE module. At load time the DIFFUSE module registers a new raw socket option which is used to configure feature, classifier and export instances, as well as for showing and deleting them using the raw socket interface. This is the same interface used by IPFW and Dummynet.

The DIFFUSE module also registers itself with the IPFW module (which must be loaded before DIFFUSE can be loaded). After DIFFUSE has registered the IPFW module will call DIFFUSE hooks every time an IPFW rule is added or removed with an action or option unknown to IPFW. This allows the DIFFUSE module to handle the instantiation and removal of new rule actions and options, such as the mlclass action or the match-if-class option.

The IPFW module also calls a DIFFUSE hook for every packet that is checked when there are rule actions and options unknown to IPFW. This allows the DIFFUSE module to process the new actions or options, and decide whether a packet matches or not.

Since the IPFW control interface (based on raw socket options) does not allow unsolicited messages from kernel to userspace and frequent polling of the kernel classifier ipfw\_col [-c actions-file] [-r min-rule-no["-"max-rule-no]] [-s sctp-port] [-t tcp-port] [-u
udp-port][-nq]

Figure 18. Collector configuration

```
actions-file = 1*line
line = ( # comment / default / classifier-name:class_number) action [action_parameters]
```

Figure 19. Actions file syntax

```
# class 0 is rt and class 1 is non-rt
default queue 2
myclass:0 queue 1
myclass:1 queue 2
```

Figure 20. Actions file example

```
ipfw_exp -c localhost -a sctp://action1.node:5000 -q
ipfw_col -c class_actions.txt -r 10000-20000 -s 5000 -t 5000 -q
```

Figure 21. Exporter and Collector configuration example

is impractical, a separate interface (UDP socket) is used to convey flow rules to the Exporter (IPFW-EXP). The DIFFUSE module only exports flow rules, if there are any rules with the new export action. The Exporter receives the flow information and forwards them to ANs, possibly using different transport protocols, such as SCTP or TCP (see Section VII).

Users use the DIFFUSE-specific config, show and delete commands as well as the new rule actions and options via an extended ipfw userspace tool (see Section V). A modified version of WEKA [17] generates classifier models based on training data. The extended ipfw userspace application parses the models and configures the DIFFUSE kernel module.

Figure 23 shows the internals of the DIFFUSE kernel module (dashed lines indicate relations between objects and solid lines indicate message flows). Feature and classifier algorithms are actually separate modules. The config commands create configured instances of these algorithms, which are kept in linked lists. Configured export instances are also stored in a linked list. DIF-FUSE actions and options in IPFW rules point to these instances. Flow information, computed features and flow classes are stored in a flow table. Flow rules are stored in a first in first out queue and later exported via the CP.

Flow information, such as 5-tuples, flow state (e.g. TCP state, timeouts), computed features and flow classes are stored in a flow table, which is realised as hash table with last recently used sorting of the bucket lists. Export

rules create flow rules that are stored in a first in first out (FIFO) list and later exported via the control protocol.

2) Flow table: The flow table stores the active bidirectional packet flows (5-tuple), their current feature statistics and assigned classes. The flow table is implemented as hash table with last recently used sorting of the bucket lists. For consistency we use the same XOR-based hash function IPFW uses for dynamic rules. This hash function is very fast to compute, and since it is commutative only one computation is required for bidirectional flows. However, depending on the flows' 5tuples it may produce a sub-optimal (non-uniform) hash value distribution. Improving the hash function is left for future work.

3) Flow timeouts: Flows are ended by configurable timeouts that depend on the protocol (UDP or TCP) and for TCP it also depend on the flows' state (connection establishment, running or teardown). If explicit rule removal messages (see Section VII) are not needed, expired flows are only freed once a new flow is inserted to the same bucket. However, if rule removal messages are required, timely flow timeouts are needed. This is implemented using a variation of a timing wheel [18], supporting one-second precision timers.

Our timing wheel is a array of double-linked lists. Each entry in the array corresponds to a second, and the entry's list holds all the timers that expire at this second. The current time is indicated by a pointer that moves through the array, wrapping around from end to start

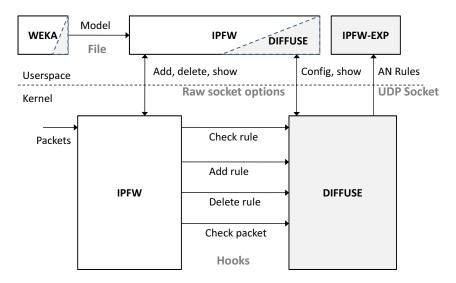


Figure 22. Classifier node design, main building blocks

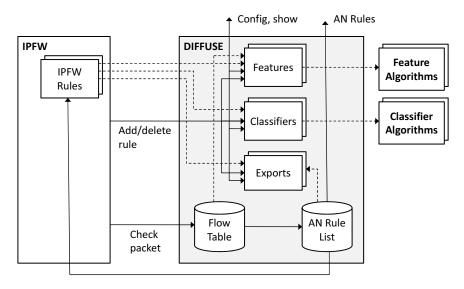


Figure 23. Classifier node design, DIFFUSE module

(circular array). This data structure allows adding and expiring of timers, as well as removal of expired timers with O(1) computational complexity. To maintain timer accuracy and avoid expiry of many timers at once, the timer wheel is checked for expired timers every 100 ms.

4) Feature computation: DIFFUSE can compute feature statistics over (overlapping) sliding windows or nonoverlapping "jumping windows". Let w be the window size in packets. Sliding windows allow more frequent feature statistic updates but have  $O(w^2)$  computational complexity and O(w) memory complexity. Jumping windows update statistics less frequently but have only O(w) computational complexity and only O(1) memory complexity in the best case (if statistics can be computed without the need to store per-packet information, e.g. to compute the mean of the packet length only the sum of the packet lengths and the number of packets is needed) or O(w) memory complexity otherwise.

By default DIFFUSE will only classify flows once the first window has been filled. To minimise classification latency, DIFFUSE also supports classification of partial windows. If enabled DIFFUSE will classify flows as soon as at least one statistic is available (e.g. for interarrival times at least two packets are needed).

5) Independent rules: IPFW/DIFFUSE rules are independent of each other. This is beneficial because it allows ruleset modifications (adding and deleting rules) on the fly. However, a consequence is that DIFFUSE must check during run-time whether the features needed by a newly added rule are already computed because of previous rules or must be added to the set of features that need to be computed, but this has only minimal performance impact.

6) Classifier sampling: To improve performance DIF-FUSE supports randomly sampled classification. Feature state is updated for every packet (e.g. the packet length is stored), but feature statistics (e.g. the mean) are computed and the classifier is executed only for sampled packets. The class of the last sampled and classified packet of a flow (if any) is assigned to non-sampled packets. However, the first sub-flow of a flow is always classified to minimise classification latency. Furthermore, rules can be configured to match only if n consecutive sub-flows were classified as the same class.

7) Passive and active CN: Classifications made by DIFFUSE can be used straight-away to decide the fate of local or remote (via ANs) packets, such as allow, block or prioritise packets (*active CN*). However, DIFFUSE can also be used for passive monitoring, such as collecting traffic statistics (*passive CN*).

8) Locking: Locks are needed to protect key data structures since they are accessed based on arriving packets as well as management from userspace (the ipfw command) and IPFW/DIFFUSE supports concurrency on multi-processor machines (nearly every recent PC). The feature, classifier and export lists are protected by a "main" read-write lock. The flow table and flow rule queue are each protected by separate read-write locks.

When a packet is inspected by IPFW/DIFFUSE the main lock is only acquired in read mode, hence packets can be processed in parallel (and processing is blocked only when rules are modified). However, during part of the processing packets may block each other, because the two other locks must be acquired in write mode when flow table or rule queue need to be modified. (Flow timeouts also block access to the flow table.)

#### B. Action Node

Figure 24 shows the main building blocks of the AN. The Collector (IPFW-COL) listens for flow rules from CNs and configures the packet filter and traffic shaper accordingly using existing configuration interface(s). We implemented our Collector as a front-end (handling the CP communication and managing addition/removal of flow rules stored in an internal database) and back-end (which generates rules customised for the underlying packet filter or traffic shaper). Our implementation creates IPFW/Dummynet rules via the command line (ipfw add). Flow information in the database is deleted based on rule removal messages and timeouts, which triggers removal of the corresponding IPFW/Dummynet rules (ipfw delete). On the AN an unmodified IPFW/Dummynet can be used (without DIFFUSE extensions).

The AN is not limited to operating systems with IPFW/Dummynet. A front-end can be written in any suitable language, and support for other firewalls or traffic shapers can easily be provided by alternative back-ends. The back-end might also implement actions such as logging of classified traffic in a database.

# C. Implementation Language(s)

The kernel part of DIFFUSE is implemented using the C programming language. To be able to share code between userspace and kernel and achieve high performance we decided to also implement the Exporter and Collector using C/C++. Since WEKA is implemented in Java, our WEKA extension is also implemented in Java.

# VII. DESIGN OF REMOTE ACTIONS PROTOCOL

We first discuss the requirements and then describe the design of the protocol.

# A. Transport Protocol

UDP and TCP are the main transport protocols currently used in IP networks. UDP is a connectionless protocol providing unreliable data transport without flow and congestion control. Due to its simplicity the overhead (both in terms of network capacity and CPU utilisation) is lower than for TCP. Furthermore, it allows more precise timing of messages by the sender.

TCP on the other hand is a connection-oriented protocol and provides flow and congestion control as well as reliable transport of data at the cost of higher overhead and less control for the sender. TCP overhead can be somewhat reduced by measures such as persistent connections (e.g. used by web browsers and servers).

SCTP is a newer transport protocol that provides a number of advanced features that are very useful for DIFFUSE [19]. SCTP is more reliable than TCP, as it has a stronger checksum, and supports transparent failover a) over different network interfaces of one host and b) over different hosts due to its multi-homing capability. SCTP allows out of order delivery of data, which prevents the head of the line blocking problem inherent in TCP. Furthermore, SCTP allows to bundle different independent data transfers (called streams) into

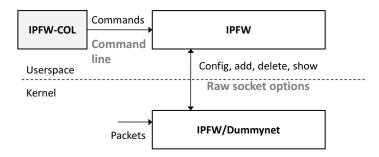


Figure 24. Action node design

one connection (called association), so only a single socket is needed. PR-SCTP is an extension of SCTP that also provides timely unreliable data transport (avoiding unnecessary retransmissions) with congestion control.

By default SCTP will use multiple network interfaces for one association, so it provides fail-over in case an interface on either side of an association becomes (temporarily) unusable. Furthermore, SCTP allows oneto-many socket associations, which can be used by a CN to transmit the same message to multiple ANs simultaneously. One-to-many socket associations can also be used to provide fail-over, e.g. if there are redundant ANs.

Our main criteria for the transport protocol are reliability, timeliness of message reception, congestion control and overhead (network and computational at the sender/receiver). We now discuss our requirements on the transport protocol taking into account different scenarios in which DIFFUSE could be used.

For traffic prioritisation one might not need maximum reliability, but it depends on the business case. If customers pay money for an improved QoS then it should be a very reliable service. For security-based applications often very high reliability is required. For market research high reliability is probably not needed. Very importantly a timely message delivery is needed for traffic prioritisation or security-based applications.

In a closed network that is dimensioned properly congestion control may not be necessary. But the Internet Engineering Taskforce (IETF) mandates the use of congestion control in the Internet (as demonstrated during the standardisation of the IPFIX protocol [20]). The overhead of the transport protocol is a less important criteria, since we usually would not have extremely short messages. With SCTP reliability is tunable and inversely proportional to overhead, but even when completely unreliable SCTP still has more overhead than UDP.

Another issue is current deployment. UDP and TCP are generally supported by every end host and network

device. SCTP is generally available on all end hosts using common modern operating systems and supported by many network devices. Table III classifies UDP, TCP, and SCTP according to the criteria identified above on the scale: -- (worst), -, +, ++ (best).

We selected SCTP as default transport protocol for DIFFUSE v0.4 because it is very reliable, provides timely message delivery, with SCTP-PR reliability and overhead can be tuned, and it provides congestion control even in unreliable mode. In situations where reliability is not an issue or there is no packet loss and congestion control is not an issue (closed well dimensioned network), UDP may be used to provide a timely message delivery with minimum overhead. TCP may be used if reliability or congestion control are required and SCTP is not available (backwards compatibility). Which transport protocol is used can be controlled by the configuration of CNs and ANs.

#### B. Protocol format

The protocol is designed to have minimum overhead while still being flexible enough to allow further extension in the future. Flexibility is crucial, because although we outlined some scenarios in which DIFFUSE could be used, we think there are many other possible scenarios.

To avoid gratuitously reinventing a wheel, our protocol is conceptually based on the IP Flow Information Export (IPFIX) protocol [20], which was developed for very similar requirements [21]. Our protocol uses the same template-based approach and similar binary encoded messages.<sup>6</sup> However, the format of protocol headers and fields is not identical to IPFIX.

1) Fixed header: Every message of the protocol has a fixed header comprised of (see Figure 25):

• Version (16 bit)

 $<sup>^{6}</sup>$ A binary-based encoding benefits DIFFUSE, since it is easy and efficient to parse in C/C++ code (especially in C kernel code), and has less network overhead compared to text-based encoding.

Table III
CRITERIA OF TRANSPORT PROTOCOLS

Protocol	Reliability	Timeliness	Cong. Control	Overhead	Deployment
UDP	_	+	—	++	++
ТСР	+	_	+	_	++
PR-SCTP	- to ++	+	+	+ to	+

- Message length (16 bit)
- Sequence number (32 bit)
- Timestamp (32 bit)

Version specifies the protocol version. Message length is the total length of the message in octets including the fixed header. The sequence number numbers all messages. It is required to determine the order of messages (in case UDP or unordered SCTP is used and packets are reordered), can be used for retransmission of information over UDP and also provides weak security against insertion attacks with UDP, as packets with sequence numbers out of the acceptable window will be silently ignored.

The Timestamp contains the time the message was generated (in seconds since Unix epoch). It allows the Collector to determine when a message was sent by the Exporter, i.e. how old the information is (assuming clocks are roughly synchronised). The collector can use this information to adjust timeouts or ignore outdated information.

2) *Templates and data sets:* After the fixed header each message contains a number of sets. Currently there are three types of sets:

- Options template
- Flow rule template
- Option and flow rule data

Flow rule data sets are used to transmit flow rules. Option data sets are used to transmit optional data. Optional data can be transmitted with different frequencies, e.g. on a per connection/association basis or on a per message basis. Options and flow templates specify the types of information elements (IEs) contained in options/flow datasets.

Each dataset has a fixed header which contains the following fields:

- Set ID (16 bit)
- Length of set (16 bit)

Set ID specifies whether the data is an options template (set ID = 0), a flow rule template (set ID = 1) or data (set ID  $\geq$  256). Length is the length of a template or data set in octets including the set header.

An options or flow rule template set contains the following fields:

- Template ID (16 bit)
- Flags/reserved (16 bit)
- A number of field definitions each consisting of an IE ID (16 bit) and optionally the length of the data in octets (16 bits)

The template ID specifies an ID for the particular template that is then referenced in a dataset (values 256– 65535). The following 16 bits are reserved for future use. Each IE ID specifies an information element, e.g. the source IP address. The length defines the length of the data in octets, e.g. length is 4 bytes for an IPV4 source address.

There are three types of IEs: fixed-length, variablelength and dynamic-length. For fixed-length IEs the IE ID also specifies the length (e.g. source IP address) and the next field is another IE ID. For variable-length IEs (e.g. a string) the length of the IE must be specified in the template in the length field following the IE ID. The length of dynamic-length IEs varies with each entry in a dataset. The first octet of a dynamic-length field in a dataset specifies the length of the field in octets (including the length field).

Whenever possible fixed-length and variable-length IEs should be used. Dynamic-length IEs should only be used if the IE length is unknown in advance and can vary significantly between entries in a data set. The highest two bits of the IE ID specify the type. If set to 00 or 01 the IE is fixed-length, if set to 10 the IE is variable length, and if set to 11 the IE is dynamic-length. This means IDs 0–32767 are for fixed-length IEs, IDs 32768–49151 are for variable-length IEs, and IDs 49152–65535 are for dynamic length IEs.

A dataset contains the data for all the IEs specified in the template in exactly the same order as specified in the template. Note that sets must aligned on 32-bit boundaries. Padding octets must be added at the end of sets as needed to ensure this.

One could reduce the overhead of the protocol by using 8-bit integers instead of 16-bit integers for IDs.

(	) 8		16	24 31		
	Version		Len	gth		
	Sequence number					
	Timestamp					

Figure 25. Fixed header of the DIFFUSE v0.4 control protocol

However, it has been shown many times that original protocol designers severely underestimated the future need for numbering space thus necessitating protocol redesign or the use of bad hacks later on. Furthermore, one could use loss-less compression to reduce the size of the data on the wire, e.g. adaptive Huffman coding is successfully used by First-person Shooter Games to reduce message sizes. However, compression increases the complexity and requires extra computational resources at the sender/receiver.

3) *Template management:* Depending on the transport protocol there are two ways of handling templates:

- Transmit templates in every packet (UDP)
- Transmit templates only at the start of connections/associations (SCTP, TCP)

If UDP is used, templates are transmitted in each packet. Packets are self-contained and loss of packets with templates will never cause loss of data beyond the lost packet (data that cannot be used because the template is not known). No additional reliability is needed for templates and there are no issues if an AN restarts. Also, an AN does not need to store templates. However, network protocol overhead is higher.

If TCP is used the Exporter only needs to transmit templates once at the start of a connection, because TCP provides ordered reliable transport. The Collector must store the templates for the duration of the TCP connection. However, if a connection is closed and reestablished, the Exporter must (re)send all templates at the start of the new connection, because it cannot know if only the connection was closed or the Collector had to restart and may have lost templates. Transmitting the templates only at the start reduces the overhead, but requires ANs to store the templates.

If SCTP is used templates and data are transmitted reliably in the same way as for TCP. If SCTP-PR is used templates only need to be transmitted once at the start, but they must be transmitted reliably using ordered reliable SCTP. The receiver must store templates for the duration of an association. Data may be transmitted unreliably if reliable data transport is not needed. Templates and data must be send over two different SCTP streams, so there will be two streams per SCTP association (templates are send over stream 0 and data is send over stream 1).

4) Information elements (IEs): The following IEs are defined (the numbers in parenthesis are the ID and the size in octets or "V" for variable-length or "D" for dynamic-length IEs):

- IPv4 Source Address (1, 4)
- IPv4 Destination Address (2, 4)
- Source Port (3, 2)
- Destination Port (4, 2)
- Protocol (5, 1)
- IPv6 Source Address (6, 16)
- IPv6 Destination Address (7, 16)
- IPv4 Type of Service (ToS) (8, 1)
- IPv6 Flow Label (9, 4)
- Class Label (10, 2)
- Match Direction (11, 1)
- Message Type (12, 1)
- Timeout Type (13, 1)
- Timeout (14, 4)
- Action Flags (15, 1)
- Action (32768, V)
- Action Parameters (32769, V)
- Classifier Name (32770, V)
- Export Name(32771, V)
- Class Tags (49152, D)

A number of IEs are fields taken from the IPv4 or IPv6 headers. Other IEs are explained briefly in the following paragraphs. Class Label is the flow's class assigned by the classifier. Classifier Name specifies the classifier that classified the flow. Match Direction indicates whether matched flows are unidirectional or bidirectional. Message Type specifies whether a message is an "add" or "remove" message. Timeout Type specifies whether the timeout is a flow timeout or a rule timeout. Timeout is the timeout value in seconds.

Action Flags are used to indicate whether the AN should apply actions to unidirectional or bidirectional flows. Action is the action name and Action Parameters specify the parameters of the action. Export Name is the

CAIA Technical Report 110704A

name of the export that generated the flow rule. Class Tags defines a list of classifier name and class tuples (if flows were classified by multiple classifiers).

The current implementation uses a standard template with the following fields: Export Name, Message Type, IPv4 Source/Destination Address, Source/Destination Port, Protocol, Class Tags, Timeout Type, Timeout, Action Name, Action Flags and Action Parameters. With a single class tag the size of one data entry is 54 bytes. (Note that IPv6 is not fully supported yet.)

5) *Keep-alive:* To minimise the overhead of establishing and shutting down connections repeatedly, a TCP connection or SCTP association between CN and AN is kept open even if there is nothing to send for a while. This keep-alive mechanism is based on the default SCTP or TCP keep-alive mechanisms. UDP is connection-less and there is no overhead, hence for UDP no keep-alive mechanism is used.

# C. Example

Figure 26 shows an example message consisting of the fixed header, a template set defining the IEs for template 256, and a flow rule data set for template 256 with multiple rules.

#### D. Security Considerations

We assume that often CNs and ANs are part of the same trusted network, which includes being connected via a secure Virtual Private Network (VPN). This prevents alteration or eavesdropping attacks on messages in flight. However, if CNs and ANs are connected via an untrusted network the integrity of messages must be protected against attackers by using digital signatures or encryption. If messages contain sensitive information, encryption must be used to preserve message confidentiality.

An AN needs to authenticate messages; it must verify that messages were created by a trusted CN. Otherwise, attackers could send fake messages to ANs for various purposes including but not limited to obtaining services that they have not paid for (e.g. prioritisation of traffic) or mounting Denial of Service (DoS) attacks by blocking a victim's legitimate flows.

A simple message authentication could be based on IP addresses. An AN only accepts messages from specified CN IP addresses on the specified ports using the specified protocols. If IP spoofing can be prevented the system will be reasonably secure in many cases. If IP spoofing is possible the sequence numbers provide some protection against blind insertion attacks. However, if strong protection against such attacks is required cryptographic authentication of messages must be used.

Like IPFIX strong security for our protocol can be provided by the Transport Layer Security (TLS) [22] or Datagram Transport Layer Security (DTLS) [23] protocols. For UDP and PR-SCTP DTLS must be used. For TCP TLS must be used.

The current implementation does not support message authentication or confidentiality.

#### VIII. MODEL CREATION & OFFLINE ANALYSIS

Since version 0.3 DIFFUSE also includes functionality to perform offline analysis based on trace files. Most importantly this allows extracting the feature statistics from traffic traces that subsequently can be use for training classification models. It also allows analysing the classification performance depending on various parameters without time-consuming online experiments. Often traffic datasets may cover many hours, and since online experiments (e.g. replaying traffic across a real network) have to be carried out in real-time, the resulting experimental time would be enormous.<sup>7</sup>

Furthermore, DIFFUSE's offline capabilities enable developers to test the functionality and correctness of feature and classifier module implementations in userspace. Testing new modules in userspace is much more convenient than testing them inside the kernel. Also, for verification the output of DIFFUSE classifiers can be directly compared to the output of ML algorithms supported by WEKA.

DIFFUSE provides two tools for offline work. The command ipfw\_fstats allows extracting sub-flow feature statistics from existing traffic traces in tcpdump format.<sup>8</sup> It uses the same feature modules and flow table implementation as used for online classification. The command ipfw\_cltest uses a DIFFUSE model to classify sub-flow instances provided as WEKA ARFF file. It uses the same classifier modules implementation as used for online classification.

Figure 27 shows how to use  $ipfw_fstats$ . The -d switch specifies a tcpdump file. The -f switch specifies a comma-separated list of feature statistics (the same names that are used for the commands described in

<sup>&</sup>lt;sup>7</sup>Even online experiments with only packet-length features must be done in real-time, because packets are grouped into flows using timeouts. Speeding up the traffic replay changes the flow structure and alters results.

<sup>&</sup>lt;sup>8</sup>Currently, only the tcpdump format is supported, but there exist tools to convert other types of trace files into tcpdump format. For example, trace files in Endace Record Format (ERF) can easily be converted to tcpdump format.

0	8	16	24	31
	Version Length			
Sequence number				
	Timestamp			
	Set ID = 1 Set length			
	Template ID = 256		Flags	
	ID = Src IP		ID = Dst IP	
ID = Classifier ID = Class				
	Set ID = 256		Set length	
	192.168.0.1		192.168.0.2	
	myclass		1	
	192.168.0.3		192.168.0.2	
myclass 2				

Figure 26. Example of DIFFUSE v0.4 control protocol message including a flow rule template and dataset

Section V). The -F option allows specifying tcpdump filter strings (see man tcpdump); only traffic in the trace file matching the filter string will be processed. By default  $ipfw_fstats$  uses its own output format, but specifying -W will generate WEKA ARFF files. The -q switch silences some additional output.

The  $-\circ$  switch specifies a feature options file that contains the feature configuration. The feature options file is a text file with one line for each feature. Each line starts with the feature module name followed by a list of options specified in the same way as for a feature config commands (see Section V-C). Figure 28 shows a small example feature options file.

The output of  $ipfw_fstats$  is a header followed by a table. Each row in the table contains the following columns: flow ID, number of packets of the flow thus far (starting with zero), protocol number (e.g. 6=TCP), source IP, source port, destination IP, destination port, and the feature statistics specified in the order they were specified with -f. The header contains a definition of all columns. The option -c will cause the specified class name to be printed in the last column. Figure 29 shows a small example of the output (with using -c or -W). (If the option -n is used the output of the header is suppressed.)

Instead of printing out all existing sub-flows the -p, -P and -S switches make it possible to sample the data. Either -S p is used to specify periodic sampling and -P is used to specify the period, or -S r is used to specify random sampling and -p is used to specify the probability. (The -s switch can be used to specify the random generator seed.)

The -r switch will cause ipfw\_fstats to print a second row for each (sampled) sub-flow, where the statistics in forward and backward directions are reversed. Effectively, this implements the Synthetic Sub-flow Pairs (SSP) approach described in [24]. For unidirectional features ipfw\_fstats can do the reversal without further input, but for bidirectional features the feature statistics to be reversed must be explicitly specified with the -R switch. This switch specifies a commaseparated list of feature indices (starting with zero) of the additionally-generated feature-reversed sub-flow instance. For example, if there are two feature statistics, the first in forward direction and the second in backward direction, the command line option -R 1, 0 means both are reversed in the feature-reversed instance.

Figure 30 shows how to use  $ipfw\_cltest$ . The -a parameter defines the algorithm name, which is the name of the classifier module to be used (see Section V-F). The -c and -f options defines the index of the class name and the indices of the feature statistics to be used respectively (starting with zero). The indices refer to the attribute number in the test file specified with -T. The test file must be in WEKA ARFF format. The -m switch specifies the classifier model created with the modified WEKA. The -q switch silences additional output. The output of  $ipfw\_cltest$  is in the same

```
ipfw_fstats -d dump-file -f feature-stat,... [-c class-name] [-F pcap-filter] [-o
options-file] [-p sample-probability] [-P sample-period] [-R reverse-idx-list] [-s seed]
[-S sample-strategy] [-nqrW]
```

Figure 27. DIFFUSE v0.4 feature statistics extractor (ipfw\_fstats)

plen window 15 payload-len iat window 20 accurate-time

Figure 28. Example feature options file

# flow-id, pkt-cnt, proto, src-ip, src-port, dst-ip, dst-port, fwd.min.plen, fwd.mean.plen, fwd.max.plen 0, 0, 6, 10.0.0.1, 22, 10.0.0.2, 43282, 48, 48, 48 0, 1, 6, 10.0.0.1, 22, 10.0.0.2, 43282, 48, 48, 48 [...]

Figure 29. Example ipfw\_fstats output

format as the output produced by WEKA, when used from the command line to test an ARFF file with an existing classifier model.

Figure 31 shows an example on how to use both commands. It shows how to extract the sub-flow statistics, build a DIFFUSE model and test the model with the training data.<sup>9</sup>

#### IX. SELECTED ML TECHNIQUES

We leverage WEKA [17] to perform the initial training data analysis and to build classifier models used for classification. WEKA provides an easy to use GUI as well as a command line interface to inspect the data, experiment with different classification techniques and build models from training data. After a classifier model has been trained in WEKA it can be saved and used with DIFFUSE v0.4.

WEKA provides access to many different ML algorithms. We previously investigated the use of different ML techniques for the classification of network traffic and found that the better techniques provide similar accuracy, but differ greatly regarding training time and classification speed [25]. Our current implementation of DIFFUSE supports C4.5 and Naïve Bayes.

We use the C4.5 decision tree classifier because it provided good accuracy for network traffic classification previously [25], the classification function is fast (tree search) and relatively easy to implement (unlike other algorithms it does not require mathematical functions not implemented in the FreeBSD kernel). Using a decision tree algorithm has the advantage that a human can interpret the resulting classifier (decision tree), although with increasing size this becomes difficult.

Naïve Bayes also was previously used to classify network traffic [25]. While the achieved accuracy of Naïve Bayes was lower than for C4.5, the classification function is fast and very easy to implement. Due its simplicity Naïve Bayes is significantly quicker than C4.5 in building a classifier (training).

More ML algorithms could be supported in the future. However, it may be difficult to implement some ML classifiers in kernel space because of the lack of mathematical functions. Future work could also investigate ways of diverting packets (or a list of features) to a userspace application that performs the classification.

### A. C4.5

C4.5 creates a classifier based on a tree structure of nodes, branches and leaves [26]. Nodes in the tree represent features, and branches represent value tests. A series of nodes and branches is terminated by a leaf, which represents the class. Determining the class of an instance is simply a matter of tracing the path of nodes and branches to a terminating leaf node.

C4.5, as other decision tree learners, uses the 'divide and conquer' method to construct a tree from a set of training instances S. If all cases in S belong to the same class, the decision tree is a leaf labelled with that class. Otherwise the algorithm will use tests to divide S into several non-trivial partitions.

<sup>&</sup>lt;sup>9</sup>For training the model with WEKA we need to generate an ARFF file with only the feature statistics. Note that for WEKA feature indices start at *one*, but for <code>ipfw\_cltest</code> feature indices start at *zero*.

ipfw\_cltest -a algorithm-name -c class-index -f feature-index,... -m model-file -T test-file
[-q]

Figure 30. DIFFUSE classifier tester (ipfw\_cltest)

ipfw\_fstats -d traf1.dmp -f fwd.min.plen,fwd.max.plen -o fopt.txt -c class1 -qW > train.arff ipfw\_fstats -d traf2.dmp -f fwd.min.plen,fwd.max.plen -o fopt.txt -c class2 -nqW >> train.arff java weka.filters.unsupervised.attribute.Remove -i train.arff -R 1,2,3,4,5,6,7 -o train\_r.arff java weka.classifiers.trees.J48 -B -t train\_r.arff -y > model.c45.diffuse ipfw\_cltest -a c4.5 -c 9 -f 7,8 -m model.c45.diffuse -T train.arff

Figure 31. Example commands for extracting sub-flow features, training a model and testing the model with the training data

Each of the partitions becomes a child node of the current node and the tests to separate S are assigned to the branches. C4.5 uses two types of tests each involving only a single attribute A. In case of discrete attributes the test is A =? with one outcome for each value of A. For real attributes the test is  $A \leq \theta$  where  $\theta$  is a constant threshold. To find the optimal partitions C4.5 relies on greedy search and selects the test set that maximizes an entropy-based gain ratio.

The divide and conquer approach partitions the data until every leaf contains instances from only one class or a further partition is not possible because two instances have the same features but different class. If there are no conflicting cases the tree will correctly classify all training instances. This over-fitting leads to a decrease of the prediction accuracy. C4.5 attempts to avoid overfitting by removing some structure from the tree after it has been built (tree pruning).

Because C4.5 selects feature tests in order of maximising the entropy-based gain ratio it is not adversely affected by unimportant or irrelevant features like other techniques. The most useful features are always used at the top of the tree and irrelevant features are ignored. Feature pre-selection is not necessary, although sometimes it still improves accuracy slightly [25].

#### B. Naïve Bayes

Naïve-Bayes is based on the Bayesian theorem [25]. It estimates the likelihood that an instance belongs to a class based on the probability that the instance belongs to the class without taking any features into account (prior probability), and the conditional probability derived for the relationships between feature values and classes (from the training data). The prior probability of a class can be computed by simply counting how many times it occurs in the training dataset. The conditional probability cannot be directly computed, but under the assumption

that the features are independent it becomes the product of the probabilities of each single feature.

The Bayes formula is only applicable if all features are qualitative (nominal). A qualitative feature takes a small number of values. Then the probabilities can be estimated from the frequencies of the instances in the training data set. Quantitative features can have a large number of values (possible infinite) and the probability cannot be estimated from the frequency distribution. Instead these features must be modelled by some continuous probability distribution (often the Gaussian distribution is used). An alternative approach is to use discretisation, which transforms quantitative features into qualitative features.

Since the true density is usually unknown for realworld data, unsafe assumptions often occur when using continuous probability density functions. Discretisation circumvents this problem. On the other hand kernel density estimation can be used instead of simple density functions to model complex distributions.

In theory, a Naïve-Bayes prediction will only be correct if all the features are statistically independent of each other and the quantitative features behave according to the probability density models. However, in practice the algorithm often produces good results even when these assumptions are violated [25].

# C. Classifier Model File Format

WEKA saves classification models produced during the training as Java serialised objects. This format is relatively complicated and no reliable free open-source C/C++ parsers exist. To use a model generated with WEKA we have extended WEKA. A command line switch (-y) was added that saves WEKA models in an ASCII format, that is easily readable for DIFFUSE.

A new interface class Diffusable was added to WEKA that needs to be used by all classifiers supported by DIF-FUSE. For each classifier using the Diffusable interface we have implemented functions to export the classifier model in ASCII format. We now describe the export format using the same syntax as in Section V.

As shown in Figure 32 each model file first lists the class names and feature/attribute names. The lines following these lists are classifier algorithm specific. (Spaces and newlines are explicitly indicated here by SP and NL respectively.)

For C4.5 each line in the model file represents a tree node and the associated test. The parameter node specifies the name of the node (always  $n_X$ ) and the parameter feature specifies the name of the feature/attribute (always  $a_X$ ), where X corresponds to the numeric index of a feature in the feature list or the number of the node in the tree (starting with zero). The next parameter specifies if the feature is nominal or real. The parameter missing-class specifies the resulting class (always  $c_X$ ) if the feature is undefined (missing), where X corresponds to the numeric index of a class in the class list (starting with zero). Then the feature test is specified. There are three different cases:

- Nominal features with non-binary splits: a list of pairs of values and class/node names. Each value specifies a feature value and is followed by either a class name or node name.
- Nominal features with binary splits: a value followed by the class name or node name for equal feature values and the class name or node name for non-equal feature values.
- Real features: a real split value followed by the class or node name for lesser equal feature values and the class or node name for greater feature values.

For Naïve-Bayes the first line defines the prior probabilities of each class. The following lines define the conditional probabilities for feature value intervals (discretised features) or the parameters of the Normal distribution (non-discretised features). Naïve-Bayes with kernel-density estimation is currently not supported.

For nominal features or discretised features there is one line for each feature value (or feature value range). For each real feature there are four lines specifying the mean and standard deviation of the Normal distribution, and the weight sum and precision of the feature for each class.

As for C4.5 the parameter feature is the feature/attribute name (always a\_X), where X corresponds to the numeric index of a feature in the feature list (starting with zero). The parameters class-prior-prob and class-cond-prob are the prior probabilities of classes and the conditional probabilities of classes depending on the feature values. The class-cond-value are the class values for the different attributes of real features.

Figure 33 depicts an example of a C4.5 classifier model generated for WEKA's iris dataset [17]. Figure 34 shows the first part of a classifier model generated by Naïve-Bayes for the same dataset [17].

## X. CONCLUSIONS AND FUTURE WORK

This report presented the DIFFUSE v0.4 system, an extension for the IPFW packet filter and shaper [5] that provides ML-based traffic classification based on statistical properties and decouples flow classification and treatment (distributed firewalling). We described the basic architecture and outlined the design of the software. We also defined and explained the main interfaces of the system: the extended ruleset language, the control protocol, and the format of classifier model files.

This report is not a manual. Man pages and HOW-TOs are provided as part of the DIFFUSE v0.4 open source software release, which can be obtained from http://caia.swin.edu.au/urp/diffuse.

In future work we will analyse the system's classification accuracy, performance and scalability. We also will explore whether automatic (re)training of classifiers may be practically achieved using live IP traffic going past particular points inside an ISP network, and the degree to which noise (packet loss and jitter) in the live traffic negatively impacts on the system's ability to recognise the same class of traffic in the future.

#### Acknowledgements

This project has been made possible in part by a grant from the Cisco University Research Program Fund at Community Foundation Silicon Valley for a project titled "Exploring the efficacy of distributed statistical traffic classification using modified open source packet filters".

#### REFERENCES

- [1] S. Zander, G. Armitage, "Design of DIFFUSE v0.1 DIstributed Firewall and Flow-shaper Using Statistical Evidence," Tech. Rep. 101223A, CAIA Technical Report, December 2010. http://caia.swin.edu.au/reports/101223A/ CAIA-TR-101223A.pdf.
- [2] T. Nguyen, G. Armitage, "A Survey of Techniques for Internet Traffic Classification using Machine Learning," *IEEE Communications Surveys & Tutorials*, vol. 10, no. 4, pp. 56–76, 2008.
- [3] P. Branch, "Lawful Interception of the Internet," *The International Journal of Emerging Technologies and Society*, Spring 2003.
- [4] J. But, N. Williams, S. Zander, L. Stewart, G. Armitage, "Automated Network Games Enhancment Layer - A Proposed Architecture," in *Proceedings of 5th Workshop on Network & System Support for Games (NetGames) 2006*, October 2006.

CAIA Technical Report 110704A

classes 1\*(class-name SP) NL
attributes 1\*(attribute-name SP) NL
( c45-model / nbayes-model )
c45-model = 1\*( node feature (n / r) missing-class 1\*(value (class / node) SP) /
1\*(split-value (le\_class / le\_node) (gt\_class / gt\_node) SP) / 1\*(value (match\_class /
match\_node) (no\_match\_class / no\_match\_node) SP) NL)
phayea\_model = prior 1+(alagg\_prior\_pred SP) NL 1+(feature (feature value 1+(alagg\_prior\_pred SP)

nbayes-model = prior 1\*(class-prior-prob SP) NL 1\*(feature (feature-value 1\*(class-cond-prob SP)) / ((mean / stddev / weightsum / precision) 1\*(class-cond-value SP)) NL)

#### Figure 32. Model file format

classes Iris-setosa Iris-versicolor Iris-virginica attributes sepallength sepalwidth petallength petalwidth n\_0 a\_3 r c\_0 0.6 c\_0 n\_1 n\_1 a\_3 r c\_1 1.7 n\_2 c\_2 n\_2 a\_2 r c\_1 4.9 c\_1 n\_3 n\_3 a\_3 r c\_2 1.5 c\_2 c\_1

Figure 33. Example C4.5 model

classes Iris-setosa Iris-versicolor Iris-virginica attributes sepallength sepalwidth petallength petalwidth prior 0.33333 0.33333 0.33333 a\_0 -inf-5.55 0.90566 0.22642 0.03774 a\_0 5.55-6.15 0.07547 0.45283 0.20755 a\_0 6.15-inf 0.01887 0.32075 0.75472 a\_1 -inf-2.95 0.0566 0.66038 0.41509 a\_1 2.95-3.35 0.35849 0.30189 0.4717 a\_1 3.35-inf 0.58491 0.03774 0.11321 [...]



- [5] The FreeBSD Documentation Project, "FreeBSD Handbook, Section 30.6 IPFW." http://www.freebsd.org/doc/en/books/ handbook/firewalls-ipfw.html.
- [6] The OpenBSD Project, "PF: The OpenBSD Packet Filter." http: //www.openbsd.org/faq/pf/.
- [7] The netfilter.org Project, "Netfilter Firewalling, NAT and Packet Mangling for Linux." http://www.netfilter.org/.
- [8] S. Zander, G. Armitage, "DIstributed Firewall and Flow-shaper Using Statistical Evidence (DIFFUSE)." http://caia.swin.edu. au/urp/diffuse/.
- [9] T.T.T. Nguyen, G. Armitage, "Training on Multiple Sub-flows to Optimise the Use of Machine Learning Classifiers in Realworld IP Networks," in *Proceedings of 31st IEEE Conference* on Local Computer Networks, November 2008.
- [10] D. Reed, "IP Filter." http://coombs.anu.edu.au/ipfilter/.
- [11] P. Dibowitz, "IPF FAQ." http://www.phildev.net/ipf/index.html.
- [12] L. Rizzo, "Dummynet." http://www.iet.unipi.it/~luigi/ip\_ dummynet/.
- [13] The FreeBSD Documentation Project, "FreeBSD Handbook, Section 30.4.6 Enabling ALTQ." http://www.freebsd.org/doc/ en/books/handbook/firewalls-pf.html.
- [14] D. Hartmeier, "Design and Performance of the OpenBSD Stateful Packet Filter (pf)," 2002. http://www.benzedrine.cx/ pf-paper.html.

- [15] M. Adamo, M. Tablò, "Linux vs OpenBSD A Firewall Performance Test," ;LOGIN:, vol. 30, pp. 35–42, December 2005.
- [16] D. Crocker, Ed., and P. Overell, "Augmented BNF for Syntax Specifications: ABNF," RFC 5234, IETF, January 2008. http: //www.ietf.org/rfc/rfc5234.txt.
- [17] I. H. Witten, Eibe Frank, "Data Mining: Practical Machine Learning Tools and Techniques – 2nd Edition. Morgan Kaufmann, San Francisco, 2005.
- [18] G. Varghese, A. Lauck, "Hashed and Hierarchical Timing Wheels: Efficient Data Structures for Implementing a Timer Facility," *IEEE/ACM Transactions on Networking*, vol. 5, pp. 824– 834, December 1997.
- [19] R. Stewart, Ed., "Stream Control Transmission Protocol," RFC 4960, IETF, September 2007. http://www.ietf.org/rfc/rfc4960. txt.
- [20] B. Claise, Ed., "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information," RFC 5101, IETF, January 2008. http://www.ietf. org/rfc/rfc5101.txt.
- [21] J. Quittek, T. Zseby, B. Claise, and S. Zander, "Requirements for IP Flow Information Export (IPFIX)," RFC 3917, IETF, Oct. 2004. http://www.ietf.org/rfc/rfc3917.txt.

CAIA Technical Report 110704A

July 2011

- [22] T. Dierks and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2," RFC 5246, IETF, August 2008. http: //www.ietf.org/rfc/rfc5246.txt.
- [23] E. Rescorla and N. Modadugu, "Datagram Transport Layer Security," RFC 4347, IETF, August 2006. http://www.ietf.org/ rfc/rfc4347.txt.
- [24] T. Nguyen and G. Armitage, "Synthetic sub-flow pairs for timely and stable IP traffic identification," in *Proceedings of Australian Telecommunication Networks and Application Conference*, December 2006.
- [25] N. Williams, S. Zander, G. Armitage, "A Preliminary Performance Comparison of Five Machine Learning Algorithms for Practical IP Traffic Flow Classification," *SIGCOMM Computer Communication Review*, vol. 36, October 2006.
- [26] R. Kohavi, J. R. Quinlan, *Decision-tree Discovery*, ch. 16.1.3, pp. 267–276. Oxford University Press, 2002.

## APPENDIX A: RELEASE FEATURES

This appendix lists the main software features that appeared with each minor/major release (for more details see ChangeLog.txt in the software distribution). DIF-FUSE versions are numbered X.Y.Z, where X is the major version number, Y is the minor version number, and Z is the revision number (Z=0 for new major/minor versions is omitted, hence 0.4.0 is called 0.4). No new software features are added between different revisions, so for example version 0.2.1 is functionally equivalent to version 0.2. Version 0.1 (first public release)

- DIFFUSE system: CN (kernel module, ipfw\_exp) and AN (ipfw\_col)
- Initial feature modules: plen, iat, pcnt
- Initial classifier modules: c4.5, nbayes
- Initial classifier models: et\_vs\_other
- Initial offline analysis tools: ipfw\_cltest

# Version 0.2

- Added feature modules: skype
- Added classifier models: et\_vs\_other(updated), fps\_vs\_other, skype
- Added offline analysis tools: ipfw\_fstats, script for building models

# Version 0.3

- Added feature modules: plenbd, iatbd
- Added offline analysis tools: script for offline training and testing (based on ipfw\_fstats and ipfw\_cltest)

## Version 0.4

• Added a number of features to offline analysis tools, and a faster variant of the script for offline training and testing where testing data is piped directly from ipfw\_fstats into ipfw\_cltest