Denial of service attack detection based on a non Gaussian and multiresolution traffic modeling

P. Borgnat[(1)], P. Abry[(1)], G. Dewaele[(1)], N. Larrieu[(2)], P. Owezarski[(2)], Y. Zhang[(2)], Y. Labit[(2)], J. Aussibal[(3)], L. Gallon[(3)], A. Scherrer[(4)], L. Bernaille[(5)], K. Boudaoud[(6)],

[(1)]Physics Lab., ENS Lyon, UMR CNRS, [(2)]LAAS, CNRS, [(3)]IUT Mont de Marsan;

[(4)]LIP, ENS Lyon, UMR CNRS-INRI,A [(5)]LIP6, Paris VI, UMR CNRS, [(6)] I3S, Sophia Antipolis, UMR CNRS

Abstract

We design Distributed Denial of Service (DDoS) detection procedures based on a non Gaussian modeling of the marginal distributions of aggregated Internet traffic. The theoretical and practical relevances of this modeling is illustrated and discussed. From this modeling, various statistical distances (Mean Quadratic Distance of Kullback Divergence) between an observation and a reference time window are computed. We show and illustrate that anomalously large values observed on these distances betray major changes in the statistics of Internet times series and correspond to the occurrences of illegitimate anomalies such as DDoS attacks. Hence, thresholding these distances enables the design of attack detection procedures. Their central feature lies in their being multiresolution in nature: time series aggregated at several levels are jointly analyzed. The assessment of the statistical performance of detection procedures in Internet is a difficult issue as no repository of traffic containing well-documented attacks is available. To overcome this, we decided and chose to perform our own collection of DDoS attacks (with precisely controlled characteristics) and collected the corresponding traffic. This enables us to evaluate the performance (detection versus false alarm probabilities) of the proposed detection procedures and to show that they present satisfactory performance with a 1 min reaction time, even for attacks whose intensity is low.
Key Words: Attacks, Denial of Service, non Gaussian modelling, detection, Kullback divergence

1 Motivation

Because of its becoming the major universal communication infrastructure, Internet is also subject to attacks in growing numbers and varieties. Moreover, its aiming at providing multiple guaranteed services increases its vulnerability to their impacts. Notably, Denial of Service (DoS) have been widely used in a recent past [11]. They consist of highly damageable attacks able to degrade the network quality of service (QoS) in an hardly predictible manner. Often, this implies significant financial losses as applications requiring stable and guaranteed QoS (such as voice over IP) are increasingly used in Internet. Therefore, detecting such attacks constitute a major and challenging issue. Moreover, DoS attacks are nowadays distributed so that one cannot easily detect them. Distributed DoS (DDoS) attacks are usually generated using botnets (i.e., machines that are controlled by hackers) that collaborate together so that each of them send only a small part of the attacking traffic, which is difficult to detect close to the generating sources. Conversely, close to the target, traffic changes become huge and IDSs (Intrusion Detection System) easily detect the attack... But it is too late: network QoS is degraded and, therefore, the attack a success. Defense against DDoS in a realistic world can be based on detecting machines hacked by pirates and on applying drastic security policies to the specific traffics issued from these identified botnets. Detecting botnets can, for example, be performed using high interaction honeypots, as *Nepenthes* [4]. Such

a priori detections may help to tune IDSs security parameters in a context dependent manner: thresholds for attack detection procedures (such as those proposed here) can be set very low for potentially aggressive traffic (such as the one produced by botnet identified machines) so as to make sure that no attack is missed (low false negative rate). Conversely, in a peaceful context, thresholds can be set at higher levels, as attack are less probable, so has to decrease the false alarm (false positive) rate. Botnet identification complements IDSs development. Current IDSs have poor performance against DDoS, specially against those of low intensity, mostly because they are based on the use of attack signatures [13] or traffic profiles built on too elementary statistics (such as sample mean or standard deviation). The naturally large variability commonly observed on Internet traffic [12] is responsible for wide fluctuations of these statistics, producing untimely threshold exceeds and hence high rates of false positives, or, even worse, false negatives (cf., for instance, [11, 3, 5, 8, 16]).

A more recent set of works proposed to take into account richer forms (such as correlations or spectrum,...) of the statistical structure of the traffic to improve IDS performance (cf. e.g., [17, 10, 6, 1, 9]). The contribution proposed here follows the same line. Indeed, it is based on a modeling of the traffic marginal distributions using non Gaussian laws: the gamma laws, $\Gamma_{\alpha,\beta}$. The originality of this approach lies in its *multiresolution* nature (several aggregation levels Δ are jointly analyzed), which provides us with robust statistics (the evolution of the parameters α and β with respect to Δ), accurately taking into account the (short time) correlation structure of the aggregated traffic. This modeling is described in Section 3. The principle of the proposed detections consists of tracking changes along time in the dependencies of α or β with respect to Δ. This is achieved by computing *distances* [2] between the statistics estimated on a current time window and those obtained from an a priori chosen reference. Then, distances are thresholded to yield detections. These procedures are detailed in Section 4. A central difficulty in validating anomaly detection procedure lies in the assessment of their statistical performance. Indeed, hackers, when generating attacks, barely inform a priori their victims. Therefore, it is difficult to have at disposal traces containing attacks of a labeled and documented set of attacks that could be used to benchmark detection procedures. Therefore, we chose and decided to perform ourselves a set of DDoS attacks, whose characteristics and parameters can be varied in a **controlled** and **reproducible** manner. Attacks are described in section 2. From this database, we have been able to analyze the statistical performance (detection vs. false alarm probabilities) of the proposed detection procedures (cf. Section 4). Though artificial or over-simplified it may look, this reference database production methodology appears to us as a mandatory step for reliable developments and validation of attack detection mechanisms. Aggregated time series of our documented attack database are available upon request to other research groups willing to benchmark their own procedures, and could enter a large research-oriented data repository.

2 Experimental DDoS Attacks

Experimental Setting. We performed UDP flooding DDoS attacks using either IPERF [7] or Trinoo [15] (on computers with Linux distribution) to generate UDP flows with different throughputs. Trinoo, using a "daemon" installed on each attacking site (4 French laboratories: Mont-de-Marsan, Lyon, Nice, Paris – all partners with the LAAS of the MetroSec project, http://www.laas.fr/METROSEC), enabled us to create more complex and realistic attack profiles. The single computer target was located at LAAS in Toulouse. The traffic related to these attacks was transported via the French national network for education and research (RENATER). LAAS is connected to RENATER with an 100 Mbps Ethernet link that has not been overflowed during attacks.
The Database. The attacks were performed in a controlled way to be able to modify their characteristics (duration, DoS flow intensity, packets length and sending rate) so as to test our detection procedures on different scenarii. The configuration set is described in Table 1. In each case, traffic was collected by ourselves (for a duration from 60 to 90 minutes, the attack mostly occurring during the second-third) before, during and after the DDoS, so that regular

	t_i	T(s)	t_a	T_A(s)	D	V	I (%)
I	09:54	5400	10:22	1800	0.25	1500	17.06
II	14:00	5400	14:29	1800	0.5	1500	14.83
III	16:00	5400	16:29	1800	0.75	1500	21.51
IV	10:09	5400	10:16	2500	1.0	1500	33.29
V	10:00	5400	10:28	1800	1.25	1500	39.26
A	14:00	5400	14:28	1800	1	1000	34.94
B	16:00	5400	16:28	1800	1	500	40.39
C	10:03	5400	10:28	1800	1	250	36.93
X	14:00	5400	14:28	1800	5	1500	58.02
tM	18:21	5400	18:58	601	0.1	300	4.64
tN	18:22	3600	18:51	601	0.1	300	15.18
tT	18:22	3600	18:51	601	8	300	82.85

Table 1: **DoS Attacks.** Upper part performed with IPERF in 2005. Lower part with Trinoo in 2006. t_i, t_a, T, T_A stand respectively for the start time of the Trace and Attack and their durations (in seconds). D, V and I refer respectively to the controlled throughput of each attacking source (in Mbps), the length of each attack packet (in bytes), and the attack relative intensity (i.e., the ratio between the sum of all attack flows and the average throughput on the LAAS link during attack).

traffic can be analyzed before and after each attack. The impact of the attacks on the global throughput of the monitored link is highly variable depending on the attack parameters, going from a major impact on global traffic profile (III, IV, V and X) to attacks that are completely hidden in the global traffic. The goal of the detection procedures is to detect all kind of attacks, including those with lowest intensities, before they have a negative impact on the network QoS.

3 Traffic marginals modeling

The analyses and detections proposed in the present contribution are based on the modeling of aggregated traffic time series $X_\Delta(k), k \in \mathbb{Z}$, consisting of the number of packets observed within bins of size Δ. Equivalent analyses could be based on bytes aggregated traffic. Representative examples of such traffic aggregated time series are presented in Fig. 1, both with and without anomalies.

Non Gaussian Marginals. Experimental evidences reported in [14] lead us to propose to model the marginals of the aggregated traffic using non Gaussian distributions: the gamma laws, denoted by $\Gamma_{\alpha,\beta}$. A $\Gamma_{\alpha,\beta}$ random variable (RV) X is a positive

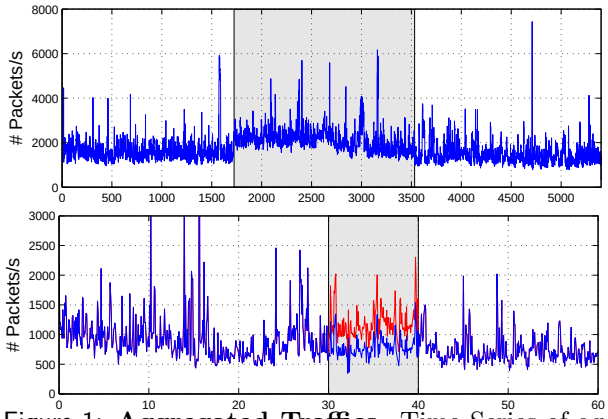

Figure 1: **Aggregated Traffics.** Time Series of aggregated traffic, $\Delta = 1$ ms, top, containing an attack (III, from 29.1 min $\leq t \leq$ 59.4 min), bottom, without attack but with an artificial anomaly (arbitrary multiplication: $X_\Delta \rightarrow \lambda X_\Delta$) 30 min $\leq t \leq$ 40 min.

RV whose probability density function (pdf) is defined as:

$$\Gamma_{\alpha,\beta}(x) = \frac{1}{\beta\Gamma(\alpha)} \left(\frac{x}{\beta}\right)^{\alpha-1} \exp\left(-\frac{x}{\beta}\right), \qquad (1)$$

where $\Gamma(u)$ is the standard Gamma function. It is fully characterized via its scale $\beta > 0$ and shape $\alpha > 0$ factors. Its mean and variance are obtained as: $\mu = \alpha\beta$ and $\sigma^2 = \alpha\beta^2$. This means that varying β for a given α simply corresponds to multiplicative increase of X, or that multiplying X only changes β by the same factor (stability under multiplication). For a given β, when α is varied from close to 0 to very large, the $\Gamma_{\alpha,\beta}$ law evolves from a strongly skewed exponential shaped-like ($\alpha = 1$) distribution to a Gaussian one ($\alpha \rightarrow +\infty$). In that respect, the quantity $1/\alpha$ can be read as an index of the distance between $\Gamma_{\alpha,\beta}$ and $\mathcal{N}(\alpha\beta, \alpha\beta^2)$ laws.

Another relevant property is the stability of $\Gamma_{\alpha,\beta}$ RVs under addition. Let X and X' be two independent $\Gamma_{\alpha,\beta}$ and $\Gamma_{\alpha',\beta}$ RVs, then $X + X'$ is $\Gamma_{\alpha+\alpha',\beta}$. This is particularly interesting when related to the aggregation procedure, because traffic aggregated at a level 2Δ is $X_{2\Delta}(k) = X_\Delta(2k)+X_\Delta(2k+1)$. Assume that $X_\Delta(k)$ is relevantly described by a $\Gamma_{\alpha_\Delta,\beta_\Delta}$ distribution, it is expected that $X_{2\Delta}$ follows a $\Gamma_{\alpha_{2\Delta},\beta_{2\Delta}}$ law. This is a theoretical argument in favor of the rel-

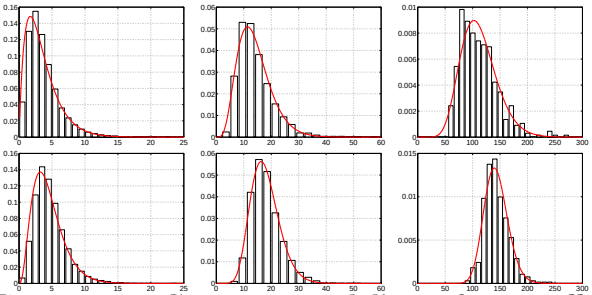

Figure 2: **Gamma modeling for traffic marginals.** Marginal distributions for aggregated traffic at three different levels $\Delta = 2^2, 2^4, 2^7$ ms (left to right) with their corresponding best fit using $\Gamma_{\alpha,\beta}$ laws. Top: regular traffic. Bottom: traffic during attack (Attack III).

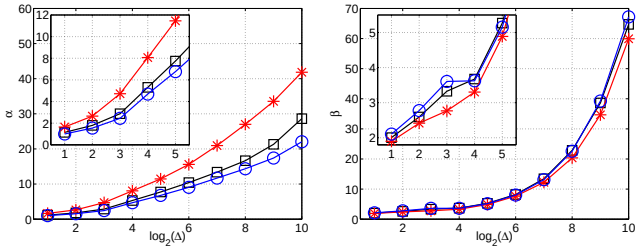

Figure 3: **Evolution of α and β with respect to the aggregation level Δ.** Left: α_Δ, right: β_Δ, averaged over adjacent non overlapping time windows ($T = 1$ min), before (blue circles), during (red asterisks) and after (black squares) the occurrence of the attack (Attack III). The evolution of α_Δ differs significantly during the attack compared to those before and after attack. Detection procedures are based on exploiting this discrepancy in the dependency with Δ.

evance of the $\Gamma_{\alpha,\beta}$ distribution. They are powerful to model the marginals of aggregated traffic for a very large range of aggregation levels, enabling a smooth and continuous evolution from stretched exponentials to Gaussians. Moreover, if the $\{X_\Delta(k), k \in \mathbb{Z}\}$ were independent RVs, one would have: $\alpha_\Delta = \alpha_0\Delta$ and $\beta_\Delta = \beta_0$ (where α_0 and β_0 are constants characterizing the packet arrival process). Any departure from these simple behaviors results from the existence of correlations amongst the $X_\Delta(k)$, and hence provide a multiresolution description of those correlations. A more elaborated version of this model [14] proposes to describe jointly the first (marginals) and second (covariances or spectra) orders of aggregated traffic: the marginals are $\Gamma_{\alpha,\beta}$ distributions while the covariances follows that of a FARIMA process, accounting jointly for short range and long range dependencies.

Goodness-of-fit. The relevance of the use of $\Gamma_{\alpha,\beta}$ distributions to model aggregated traffic marginals has been assessed by means of χ^2 goodness-of-fit tests (not reported here for sake of simplicity). We observed that the modeling is valid for Δ ranging from the millisecond up to the second (that is over three orders of magnitude). Also, we found that this modeling not only works for a regular traffic (that is for a traffic that undergoes no(-known) anomaly), but also for traffic under attack. These results hold for window size ranging from $T = 1$ min to $T = 30$ min. This is illustrated in Fig. 2 for the former choice. These finding are in agreement with those reported in [14] over a large variety of Internet traffics available from the standard major traffic repositories.

Parameter evolution with respect to the aggregation level. Let us now analyze the evolution as a function of the aggregation level Δ of the estimated shape and scale parameters. Data are split into adjacent non overlapping time window of duration T. From standard estimation techniques, referred to as mixture of *maximum likelihood* and *moment based* procedures (cf. [14]), we obtain $\hat{\alpha}_\Delta(l)$ and $\hat{\beta}_\Delta(l)$, where l refers to the time position, lT of the l-th time window. These estimates are reported in Fig. 3 (with $T = 1$ min, $\Delta = 2^1, \ldots, 2^{10}$ ms).

To analyze these results, let us use the stability under addition/aggregation property. Fig. 3, consisting of the averages (over time l and dispersions of $\hat{\alpha}_\Delta(l)$ and $\hat{\beta}_\Delta(l)$) clearly shows that the observed evolution with respect to Δ depart from the expected behavior if there were independence between $X_\Delta(k)$ and $X_\Delta(k+1)$. This is a strong evidence for the existence of significant and intricate short time statistical dependencies with in the $\{X_\Delta(k), k \in \mathbb{Z}\}$.

Furthermore, the key empirical observation lies in the fact that the evolution of $\hat{\alpha}_\Delta(l)$ and $\hat{\beta}_\Delta(l)$ functions differ notably for time windows containing regular traffic (black squares and blue circles) compared to those containing traffic under DDoS attack (red as-

terisks). Moreover, one can notice that the increase of α with Δ is much faster for traffic under attack than for regular one indicating that the convergence towards Gaussian traffic is much faster under attacks. This cannot be explained by a simple increase of the average traffic level as α is by nature not sensitive to a simple level shift (or multiplicative increase) and as the whole α_Δ function is altered. These experimental findings remarkably show that the whole structure of the short time correlations in aggregated traffic is significantly changed by the occurrence of the attack. These results betray a major change in the dynamical properties of $\{X_\Delta(k), k \in \mathbb{Z}\}$, not only in its marginal or static properties. This is the key interest of using a joint multiresolution (various Δ) analysis scheme.

4 Anomaly Detection

4.1 Detection procedures

Principles. The key ingredient leading the design of the detection procedures we propose here lies in taking advantage of the multiresolution nature of the analyses (performed jointly at different aggregation levels) and of the fact that the evolution of α and β with respect to Δ is altered when an attack occurs.

The time series under analysis are split into adjacent non overlapping time windows of length T. Independently for each time window and each aggregation level, one computes a *distance* between a statistical characteristic measured on the current time window and on a reference window. In a second step, one thresholds this distance to detect unexpectedly large deviations and hence anomalous traffic behaviors.

The design of such a detection procedure is subject to three major a priori choices: choice of the reference time window (position and duration), choice of the distance and choice of the threshold value. In the current work, we explored the following possibilities.

The reference window consists of T_{Ref} minutes of traffic collected by ourselves before the occurrence of the attacks and therefore assumed to be regular traffic. We used $T_{Ref} = 1$ min and $T_{Ref} = 10$ min.

There exists a very large variety of distances that could be used (cf. e.g., [2] for an exhaustive review). We explore here three of them chosen because they better suit the purposes and intuitions developed here and as they are also known to yield the most robust results. We compute Mean Quadratic Distances (MQD) for the functions α_Δ and β_Δ, 1D Kullback divergences for monodimensional marginal distributions of X_Δ for various Δs and 2D Kullback divergences for bidimensional marginal distributions of $(X_\Delta, X_{\Delta'})$ for various pairs $(\Delta, \Delta' \neq \Delta)$. Here, we used $\Delta = 2^1, 2^2, \ldots, 2^9, 2^{10}$ ms.

For the threshold, a collection of values is systematically explored so as to derive performance curves for the detection procedures.

Distances. To measure the distance between two (possibly multi-dimensional) probability density functions denoted by $p_1(x)$ and $p_2(x)$, one can use Kullback divergence (KD) defined as [2] :

$$KD(p1, p2) = \int (p_1 - p_2)(\ln p_1 - \ln p_2) dx. \quad (2)$$

From this definition, we compute for various aggregation levels, Kullback divergences between the marginal distribution of X_Δ estimated within the $l-$th time window, $p_{\Delta,l}$, and that obtained from the reference window $p_{\Delta,Ref}$:

$$K_\Delta^{(1D)}(l) = KD(p_{\Delta,l}, p_{\Delta,Ref}). \quad (3)$$

The $K_\Delta^{(1D)}$ at various levels are then combined together to produce multiresolution distances. To obtain directly multiresolution distances, we can also compute Kullback divergences for various pairs of aggregation levels $(\Delta, \Delta' \neq \Delta)$ between the joint two-dimensional distributions $p_{\Delta,\Delta',l}$ and $p_{\Delta,\Delta',Ref}$ obtained from the $l-$th and reference time windows, respectively:

$$K_{\Delta,\Delta'}^{(2D)}(l) = KD(p_{\Delta,\Delta',l}, p_{\Delta,\Delta',Ref}). \quad (4)$$

Second, we define the Mean Quadratic Distances (MQD) for α_Δ and β_Δ:

$$D_\alpha(l) = \frac{1}{J}\sum_{j=1}^{J} \left(\hat{\alpha}_{2^j}(l) - \hat{\alpha}_{2^j}(ref)\right)^2, \quad (5)$$

$$D_\beta(l) = \frac{1}{J}\sum_{j=1}^{J} \left(\hat{\beta}_{2^j}(l) - \hat{\beta}_{2^j}(ref)\right)^2. \quad (6)$$

Let us note that both type of distances are used in a multiresolution way, the KDs are applied directly

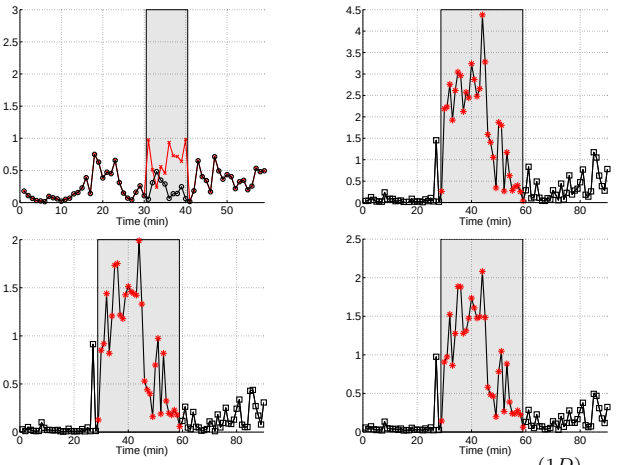

Figure 4: **Kulback Divergences.** Left, $K_{2^4}^{(1D)}(l)$ for regular traffic (top) and traffic under attack (bottom). Right, for traffic under attack, $K_{2^7}^{(1D)}(l)$ (top) and $K_{2^4,2^7}^{(2D)}(l)$ (bottom). Estimations from non overlapping 1 min time windows. While top left plot shows the natural level of the statistical fluctuations for KD, the three other plots show that the attack is clearly seen. Superimposed on top left plot $K_{2^4}^{(1D)}(l)$ for traffic containing an artificial multiplicative anomaly. Points during attacks in (red) asteriks.

to the estimated pdfs and are hence used in a non parametric way, while MQDs are explicitly applied to the functions $\hat{\alpha}_\Delta(l)$ and $\hat{\beta}_\Delta(l)$.

4.2 Distances as a response to attacks

Fig. 4 illustrates Kullback divergences as functions of the time index l, computed on the time series of regular traffic and traffic containing an attack (Attack III). Top left plot computed on regular traffic indicates the *natural* level of the statistical fluctuations of the KDs. This information can be used to set the detection threshold above this fluctuation level. The three other plots illustrate that both 1D and 2D KDs perfectly *feel* the attack and hence enable to detect it. However, because we use it in a non parametric way, KDs could be fooled by legitimate multiplicative change in traffic (The KD computed over the time series containing the multiplicative increase has been superimposed on the top left plot in Fig. 4). To

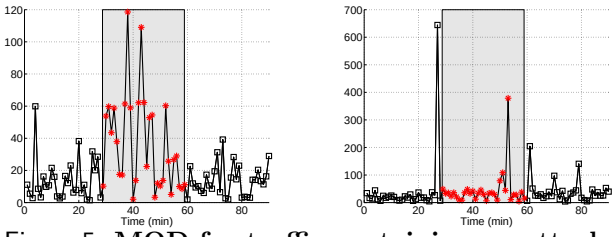

Figure 5: **MQD for traffic containing an attack.** $D_\alpha(l)$ (left) and $D_\beta(l)$ (right), computed on non overlapping 1 min time windows. Time windows containing the attack (Attack III) are shown with (red) asterisks, regular traffic with (black squares).

increase robustness, we turn to $D_\alpha(l)$ and $D_\beta(l)$.

MQDs are plotted in Fig. 5 for traffic containing an attack. The left plot in Fig. 5 clearly illustrates that $D_\alpha(l)$ takes large values within time windows l containing the attack, hence confirming that its occurrence significantly alters the dependency of α with respect to Δ. Therefore, thresholding enables detections. Conversely, one observes that $D_\beta(l)$ remains mostly stable and is not significantly shifted with the occurrence of the attack. As β is a scale parameter mostly sensitive to the intensity of the traffic, this indicates that attacks do not correspond to traffic increase but rather to significant dynamical changes in the correlation structure. The large values observed in the $D_\beta(l)$ plot correspond to time windows that do not satisfy the χ^2 goodness-of-fit test because they contain both regular and under attack traffics, yielding aberrant estimates. Note that these large values occur at the start and stop times of the attack.

A contrario, Fig. 6 presents $D_\alpha(l)$ and $D_\beta(l)$ for regular traffic. As expected, their fluctuations remain low compared to those observed on traffic under attack and would hence not lead to false alarm detections. This enables to calibrate the *natural* level for the statistical fluctuations of the distances. On top of this regular traffic, we have superimposed an artificial multiplicative anomaly which may account for a legitimate increase of traffic. By construction, $\hat{\beta}_\Delta(l)$ feels that anomaly while $\hat{\alpha}_\Delta(l)$ does not. However, the MQD $D_\alpha(l)$ and $D_\beta(l)$ do not respond to traffic multiplication as the correlation of the X_Δ and hence the dependencies in Δ of α and Δ are not altered. Therefore, together, the set of tools $\hat{\alpha}_\Delta(l)$,

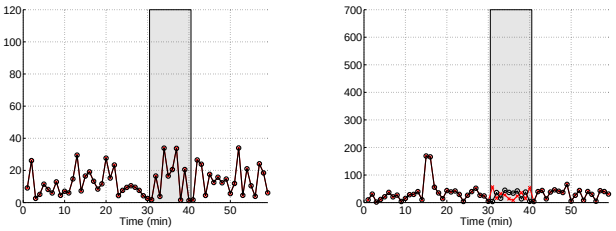

Figure 6: **MQD for normal traffic or traffic with a multiplicative increase.** $D_\alpha(l)$ (left) and $D_\beta(l)$ (right). Superimposed (and almost identical): traffic containing an artificial multiplicative anomaly (red asterisk). The MQD $D_\alpha(l)$ and $D_\beta(l)$ are not sensitive to the multiplicative increase: they stay within the size of the natural statistical fluctuations.

$\hat{\beta}_\Delta(l)$, $D_\alpha(l)$ and $D_\beta(l)$ enables us to detect anomalies in traffic statistics and classify them in legitimate ones (traffic intensity increase) and illegitimate ones (dynamical changes in the correlation structure).

4.3 Statistical performance

ROC Curves. The statistical performance of detection procedures are usually quantified via their receiver operational characteristics (often referred to as ROC curves). They consist of the plots of the correct detection vs. false alarm probabilities, $P_D = f(P_F)$, together with the plots of these two probabilities as a function of the detection threshold λ, $P_D = f(\lambda)$ and $P_F = f(\lambda)$.

ROC curves are obtained from our documented database. From the knowledge of the characteristics of the attacks, we precisely know which time windows contain the attacks and which do not. For each detection level λ, we compute the number of time windows with attacks for which the chosen distance is above threshold and derive P_D, conversely, we compute the number of time windows containing no attack and derive P_F. Then, we vary the threshold level λ by arbitrary shifts to obtain the desired functions.

Experimental results. ROC curves, obtained with the parameters $T_{Ref} = 10$ min, $T = 1$ min, $\Delta = 2^1, 2^2, \ldots, 2^9, 2^{10}$ ms, for the three different distances are reported in Fig. 7 for Attack III. These plots clearly show that the curves $P_D = f(P_F)$ are satisfactorily close to the optimal left upper corner (the ideal set point where all attacks would be detected and no false alarm raised). One sees that $K^{(2D)}$, that explicitly involves two different aggregation levels, displays better performance than a simple $K^{(1D)}$. On this example, one further notices that MQD has slightly poorer performance than those of KD. However, the former distances possess a strong robustness against legitimate anomalies and a promising ability in classifying between legitimate and illegitimate anomalies. Table 2 shows for all attacks studied here P_D for a P_F set respectively to 10% and 20%. Such tables are obtained by reading on ROC curves P_D for the a priori chosen P_F level. They show that in all cases performance are very satisfactory. Notably, for attacks with very low intensity (such as Attacks A, B, I and II) and hence little impact on traffic volume profiles, detection rates, even if low at first sight are encouraging as most traditional IDS based on simple mean and variance statistics would totally miss them.

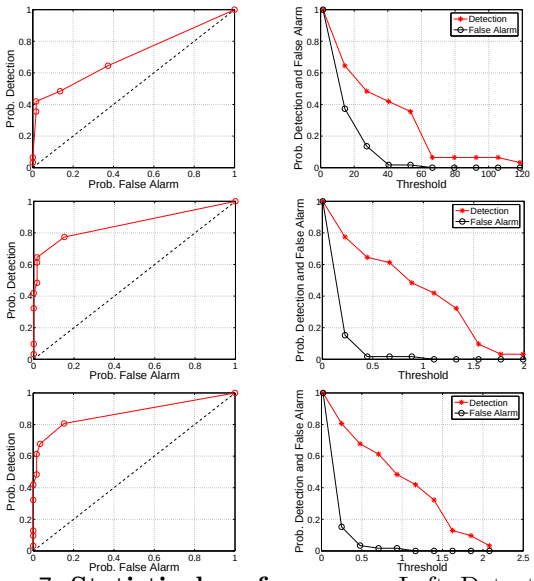

Figure 7: **Statistical performance.** Left, Detection probability vs False Alarm probability, $P_D = f(P_F)$, right, $P_D = f(\lambda)$ and $P_F = f(\lambda)$, for $D_\alpha(l)$ (top), $K^{(1D)}_{2^4}(l)$ (middle), $K^{(2D)}_{2^4,2^7}(l)$ (bottom).

5 Conclusions and perspectives

In this contribution, we designed DDoS attack detection procedures intrinsically based on a multiresolution non Gaussian modeling of traffic marginals.

	D_α	$K^{(1D)}_{2^4}$	$K^{(1D)}_{2^7}$	$K^{(2D)}_{2^4,2^7}$
I	51 : 64	25 : 64	35 : 67	25 : 51
II	48 : 54	35 : 58	35 : 61	35 : 61
III	48 : 58	74 : 93	70 : 83	87 : 93
IV	33 : 50	56 : 67	56 : 69	34 : 66
V	18 : 40	87 : 96	34 : 93	90 : 96
A	21 : 50	50 : 78	37 : 59	53 : 81
B	81 : 87	78 : 78	09 : 33	78 : 81
C	52 : 58	91 : 91	91 : 91	91 : 91
X	93 : 96	93 : 93	93 : 93	93 : 93
tM	27 : 55	36 : 91	36 : 91	45 : 91
tN	54 : 54	73 : 91	91 : 91	55 : 73
tT	82 : 82	100 : 100	100 : 100	100 : 100

Table 2: **Detection rates.** For each distance, for each attack, Detection probabilities obtained for a fixed False Alarm probability set at 10% (left) and 20% (right).

Mostly, they consist of computing (and thresholding) distances between the statistics estimated on a sliding observation time window and on a reference time window. Also, we performed a collection of DDoS attacks with controlled characteristics and collected the corresponding traffic. This enabled us to assess the statistical performance of the proposed detection procedures. We showed that results are promising as the probability detection rate remains satisfactory even for attacks with very low intensity.

To further develop the results presented here, we intend first to perform larger attack campaigns involving a larger variety of intensity levels, attack protocols or mechanisms. Second, we plan to expand on the design of detection procedures: use of other distances, detection threshold automatic selection (based on *bootstrap* techniques), further use of the multiresolution nature of our modeling. Also, a potential development lies in the choices of (peaceful or aggressive) context dependent thresholds. Third, we want to monitor the traffic simultaneously at different points in the network so as to perform joint modeling and hence collaborative detections, also it would be interesting to track changes in traffic close to the attack sources. It is expected that these different perspectives should enable to improve the performance of the procedure while decreasing the time duration of the observation window significantly below the minute.

References

[1] Barford, P., Kline, J., Plonka, D., and Ron, A. A signal analysis of network traffic anomalies. In *ACM/SIGCOMM Internet Measurement Workshop* (Marseille, France, Nov. 2002).

[2] Basseville, M. Distance measures for signal processing and pattern recognition. *Signal Processing 18* (1989), 349–369.

[3] Brutlag, J. Aberrant behavior detection in time series for network monitoring. In *USENIX System Administration Conference* (New Orleans, Dec. 2000).

[4] Dornseif, M., Holz, T., and Klein, C. Nosebreak v attacking honeynets. In *IEEE Workshop on Information Assurance and Security* (United States Military Academy, West Point, NY, USA, June 2004).

[5] Hochberg, J., Jackson, K., Stallings, C., McClary, J., DuBois, D., and Ford, J. NADIR: an automated system for detecting network intrusion and misuse. *Journal of Computer Security 12*, 3 (1993), 235–248.

[6] Hussain, A., Heidemann, J., and Papadopoulos, C. A framework for classifying denial of service attacks. In *SIGCOMM* (Karlsruhe, Germany, 2003).

[7] IPERF - The TCP/UDP Bandwidth Measurement Tool. http://dast.nlanr.net/Projects/Iperf/.

[8] Javits, and Valdes. The SRI IDES statistical anomaly detector. *ESORICS* (May 1991).

[9] Jung, J., Krishnamurthy, B., and Rabinovich, M. Flash Crowds and Denial of Service Attacks: Characterization and Implications for CDNs and Web Sites. In *International WWW Conference* (Honolulu, HI, May 2002).

[10] Lakhina, A., Crovella, M., and Diot, C. Diagnosing network-wide traffic anomalies. In *SIGCOMM* (Aug. 2004).

[11] Moore, D., Voelker, G., and Savage, S. Inferring internet denial-of-service activity. In *Usenix Security Symposium* (2001).

[12] Park, K., and Willinger, W. Self-similar network traffic: An overview. In *Self-Similar Network Traffic and Performance Evaluation*, K. Park and W. Willinger, Eds. Wiley (Interscience Division), 2000, pp. 1–38.

[13] Paxson, V. Bro: a system for detecting network intruders in real-time. *Computer Networks Journal 31*, 23–24 (1999), 2435–2463.

[14] Scherrer, A., Larrieu, N., Borgnat, P., Owezarski, P., and Abry, P. Non gaussian and long memory statistical modeling of internet traffic. In *4th Workshop IPS-MoMe* (Feb. 2006), Salzburg Research, pp. 176–185.

[15] TRINOO - Distributed network DoS tool. http://staff.washington.edu/dittrich/misc/trinoo.analysis.

[16] Vaccaro, H., and Liepins, G. Detection of anomalous computer session activity. In *IEEE Symposium on Security and Privacy* (Oakland, California, May 1989), pp. 280–289.

[17] Ye, N. A Markov chain model of temporal behavior for anomaly detection. In *Workshop on Information Assurance and Security* (West Point, NY, June 2000).