

A Wavelet-Based Framework for Proactive Detection of Network Misconfigurations

Antonio Magnaghi
Fujitsu Laboratories of America
1240 E. Arques Ave. M/S 345
Sunnyvale, CA 94085 - USA
antmagna@fla.fujitsu.com

Takeo Hamada
Fujitsu Laboratories of America
1240 E. Arques Ave. M/S 345
Sunnyvale, CA 94085 - USA
thamada@fla.fujitsu.com

Tsuneo Katsuyama
Fujitsu Laboratories Ltd.
1-1 Kamikodanaka 4-chome
Kawasaki 211-8588, Japan
katuyama@flab.fujitsu.co.jp

ABSTRACT

An increasing number of misconfigurations and malicious behaviors threaten the normal operation conditions of data networks. Thus, field engineers are constantly presented with the challenge of isolating new misconfigurations and anomalies. In this paper, we present a group of real-world problems reported by a set of six commercial networks we surveyed. Successively, we focus on a well-defined family of misconfigurations. Our analysis identifies common properties such as anomalous behaviors share. Misconfigured TCP flows experience packet losses and RTO-based (Retransmission Time-Out) events during the opening phase of the TCP connection (“Early RTO Events”). This introduces precise correlations in misconfigured traffic that we utilize as a “signature” in order to isolate the presence of anomalies. We propose a wavelet-based algorithm that is capable of revealing such a family of anomalies from the analysis of MIB data aggregating healthy and anomalous flows. Simulation and the use of real datasets from a commercial network allow us to quantitatively assess the effectiveness of our detection procedure. Numerical results show that our algorithm can effectively isolate the presence of an anomalous traffic component that is a minimal percentage of the overall link throughput. Therefore, our approach provides a general and highly sensitive misconfiguration detection instrument.

Categories and Subject Descriptors

C.2.3 [Computer-Communication Networks]: Network Operations – *network management, network monitoring.*

General Terms

Management, Measurement, Performance, Reliability.

Keywords

Misconfiguration, retransmissions, wavelets, network performance.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGCOMM'04 Workshops, Aug. 30 & Sept. 3, 2004, Portland, OR, USA.
Copyright 2004 ACM 1-58113-942-X/04/0008...\$5.00.

1. INTRODUCTION

Management and performance measurement are critical tasks in every networking infrastructure. The exceptional complexity of wide-scale networks and a social environment where malicious behaviors are routinely observed inevitably lead to a broad range of faults and misconfigurations. Such anomalies are responsible for non-optimal resource utilization and performance degradation [13, 9]. Our aim is practical in nature: we focus on the development of tools that can assist network operators and field engineers in monitoring the network “health” condition and detecting anomalous events that impact the performance of the infrastructures they operate. A detection methodology general in scope appears a necessity in order to overcome the limitations of more traditional, ad-hoc solutions, which treat each specific misconfiguration individually. One of our goals is to devise a framework capable of addressing a family of anomalies that share a common phenomenology.

The large majority (90%) of network traffic is carried by the reliable transport protocol TCP [4]. The survey we conducted with six commercial networks reveals that a significant portion of the reported anomalies is such that initial packets of the TCP three-way handshake are lost. RTO’s (Retransmission Time-Outs) follow the initial loss [14]. We define an Early RTO Event (ERE) the sequence of retransmitted packets in a state of the TCP connection such that no RTT (Round-Trip Time) measurements are available yet. ERE’s utilize default RTO values because of RTT unavailability. (We do not consider features of some O/S’s, like FreeBSD, that optionally can share historic RTT information among connections to the same destination.) RTO default values are standardized and consistently implemented in TCP/IP protocol modules. Therefore, retransmission events incurred in the opening phase of TCP connections generate network traffic with well-defined characteristics, insensitive to protocol module implementations and end-to-end path properties. ERE’s follow a deterministic pattern: the first RTO timer goes off after three seconds from the initial packet loss (effects due to clock granularity may be incurred). Successive retransmissions are regulated by an exponential back-off algorithm, until the connection is terminated. We argue that these retransmissions inside TCP flows afflicted by the misconfiguration produce exogenous dependencies. Such misconfiguration-induced correlation structures manifest properties in aggregated traffic that diverge from the typical characteristics of healthy configurations, e.g. long-range dependency and self-similarity. Based on this observation, we develop a detection algorithm capable of isolating the misconfigured component embedded in aggregated traffic. We

exploit scaling properties of the dynamics of misconfigured flows at precise aggregation levels associated with exponential back-off retransmissions of RTO events. Specifically, through wavelet analysis of the time-series of MIB packet count statistics, the energy of the input signal is decomposed at different resolution levels. Our analysis shows that ERE's are directly correlated to the presence of dips at precise resolutions. We develop a procedure to analyze the shape of the energy-based representation of the signal and infer the presence of anomalies.

The remainder of the paper is organized as follows. Section 2 summarizes the findings from our field survey and section 3 analyzes the set of specific misconfigurations we target. In section 4, a wavelet-based misconfiguration detection algorithm is presented. Successively (section 5), we assess the effectiveness of the detection algorithm in terms of simulation and analysis of real network traces. Section 6 provides further information about related work. Concluding remarks are presented in section 7.

2. SURVEY OF MISCONFIGURATIONS

Because of the practical nature of the problems under investigation, a preliminary obstacle we had to address was to characterize the problem space. This required overcoming several difficulties. In fact, commercial organizations like network operators, Internet Data Centers (IDC's) and Internet Service Providers (ISP's) regard information about anomalies afflicting their infrastructure as sensitive and confidential. We surveyed six different organizations. Each one of them periodically collects logging information of various types for monitoring, auditing and troubleshooting purposes:

- One organization is a major Tier-1 network provider, with numerous co-located access points that provide connectivity to corporate clients.
- A second organization in our survey is a nation-wide ISP.
- The remaining four entities are corporate IT departments that manage LAN's and WAN's and the services to support day-to-day business of corporations (or departments) of varying size. The smallest network in this group serves about 150 employees and the biggest one serves about 3,000 regular corporate users.

The range in size of the surveyed networks and the nature of the services provided are broad and capable of providing valuable insight. Hence, we believe that our survey represents a limited, but realistic snapshot of the problem space.

Alternative misconfiguration classification schemes are viable. For instance, we could account for the "location" of the misconfiguration (such as core network vs. end-host), or the cause of the misconfiguration (such as human error vs. HW/SW fault). However, a network-performance-centric categorization appears more suitable to our intent of detecting misconfigurations from observations of anomalies embedded in the network traffic. We, therefore, classify the anomalies in the collected logs from the standpoint of the impact they have on network performance metrics. In particular, we refer to the set of performance parameters as defined by the IETF IPPM working group [10]. We employ four metrics: packet loss, bulk transfer capacity, delay variation and packet reordering. Misconfigurations are grouped based on what metric they degrade. In cases where multiple metrics are affected by one misconfiguration, we take into account the prevailing effect of the misconfiguration. For

instance, mismatch of Medium Access Control (MAC) arbitration protocols in Ethernet-based LAN's causes certain loss patterns and consequent retransmissions. In this case, both packet loss and bulk transfer capacity are affected. However, there is a precise causality relationship between the two aspects (loss reduces data throughput). Thus, we catalog the problem in the "packet loss" family. In our field survey, relative frequencies of misconfigurations are as follows. 68% of the logged misconfigurations affects packet loss, 30% affects bulk transfer capacity. A negligible portion (2%) falls in the remaining categories. We limit our focus to packet-loss-related misconfigurations, because such anomalies appear to be dominant.

3. ANALYSIS OF TARGET ANOMALIES

From the surveyed misconfigurations, we select a group of specific anomalies to target: 1.) Duplication of IP Address Space; 2.) Packet Filtering Misconfiguration; 3.) Permanent Routing Loop; 4.) TCP-SYN Flood D-DoS attack (Distributed Denial of Service). Such anomalies account for over 33% of the cases logged in the data made available to us. Our focus is on the effects these anomalies typically have on the dynamics of the TCP transport protocol. We show how, despite the apparent discrepancies, all the target anomalies share common properties that lead us to a general detection framework (section 4).

- *Duplication of IP Address Space* is frequently observed in medium- to wide-scale networks. It is introduced when a new sub-network S-N2 is added to a network N1 or when, for maintenance reasons, the address space assigned to S-N2 is altered. Inadvertently, S-N2 address space overlaps with the address space of a different sub-network S-N1 in N1. Apparent causes can be: a.) lack of coordination among divisions administering separate portions of the same network or b.) lack of up-to-date information about recent modifications to certain network portions (incomplete network diagrams, stale configuration information...). Such a misconfiguration interferes with the internal routing state of the network. In the case a Distance Vector (DV) protocol is used, nodes in N1 close to the misconfiguration point S-N2 will change their routing state. In fact, DV information exchange reveals the existence of a shorter path to a certain prefix, namely the address space of S-N1. Once the routing state of N1 has converged, let $M(S-N2)$ be the set of routers in N1 the state of which is altered in response to such a misconfiguration. Packets addressed to S-N1 that reach a node in $M(S-N2)$ will be routed towards S-N2, where they typically are discarded. Conversely, packets addressed to S-N1 which do not reach a node in $M(S-N2)$ will be properly forwarded. Depending on the particular position inside network N1, the problem can be easily observed or completely transparent to typical monitoring activity. This increases the complexity of troubleshooting compared to misconfigurations that result in complete outages. The TCP flows affected by the misconfiguration are not able to complete the three-way handshake required to open a new connection. Different cases are possible. In Figure 1, a client in S-N1 tries to open a TCP connection to a server. S-N2 overlaps with the client address space. The server TCP-SYN-ACK packet to the client is lost in S-N2 (the shaded area is $M(S-N2)$). Hence, early RTO events take place at the client. Additionally, as the TCP-SYN-ACK from the server is lost, a half-open TCP connection is established in the server's TCP protocol module. Thus, a condition is introduced that closely resembles the TCP-

SYN flood D-DoS attack analyzed below. This, indirectly, can affect other clients (denial of service) even if their traffic does not traverse M(S-N2).




Figure 1: Possible Scenario of IP Space Duplication

- Packet Filtering Misconfiguration:** Packet filtering is a common practice in every network and it aims at improving security and integrity. In our survey, several instances of filtering misconfigurations have been reported. Generally, packet-filtering misconfigurations can result in: a.) unwanted packet drop, if the filter is excessively restrictive; b.) leaking of undesired packets if the filter configuration is too permissive. We focus on the former category of filtering misconfigurations, because the latter typically is more relevant to security issues. Excessively stringent filtering configurations can typically be attributed to several factors: a.) most supported filtering specification formats are very restrictive in their semantic, this requires administrators to write cumbersome rules; b.) filtering rules are typically packet-based, however business-centric filtering requirements are flow-oriented; c.) filtering tools impose an implicit rule-processing order that frequently is overlooked when configuration changes are made. We have identified several filtering misconfigurations that discard all packets to/from a certain address space. Such situations affect TCP connection establishment in a manner similar to the other types of target misconfigurations described in this section. The TCP handshake cannot complete and RTO-based retransmissions (ERE's) take place.

- Permanent Routing-Loop** are serious anomalies because they cause elevated bandwidth utilization and packet losses. Typically, layer-3 loops are categorized as transient or permanent [6]. Transient loops naturally occur during propagation of routing changes and disappear once convergence is reached. The logging data available to us contained instances of permanent routing-loops. Some of them were induced by erroneous static configurations of routes affecting certain prefixes. In our survey, we could also isolate a new type of persistent loop due to corruption of DV routing state. This specific anomaly appears as the interaction of plausible configuration choices in combination with misconfiguration of packet filtering. The concomitance of events is such that routing information leaks from a network N1 into an adjacent network N2. The routing state of N2 is altered in such a manner that packets sourcing from N1 are routed by N2 back to N1, typically through an interconnection point different from the one where packets from N1 entered N2 initially. Because of the peculiarity of the situation, such a misconfiguration is not frequent. However, the cases reported in our logs clearly show how detrimental it is in terms of network performance. Packets affected by the misconfiguration loop until they are eventually dropped because their TTL value expires. TCP connections initiated by hosts affected by the problem will not be able to complete and ERE retransmissions are incurred.

- D-DoS Attack - TCP-SYN Flood:** The purpose of a D-DoS attack is to harm a specific target in such a manner that the service(s) provided by the target becomes unavailable to legitimate users. Different mechanisms can be exploited [9]. We focus on the TCP-SYN flood attack, because it is a quite common practice and it causes network anomalies that manifest important analogies with the other types of misconfigurations described above. The attacker uses a set of compromised hosts from which spoofed TCP-SYN packets are generated towards the target. The target produces TCP-SYN-ACK packets destined to the spoofed addresses of the initial TCP-SYN's. TCP-SYN-ACK's from the target are, thus, lost and half-opened TCP connections saturate the incoming request queue. Subsequent incoming TCP-SYN packets are discarded when legitimate clients try to open a new connection with the target. Hence, service is denied. RTO-based retransmissions take place from the target's side (lost TCP-SYN-ACK's in response to spoofed packets) and from the clients' side (lost TCP-SYN due to overflow of queue of incoming requests at the target). The latter group of TCP flows is numerically more significant and the more successful a D-DoS attack is, the more clients' early RTO retransmissions will be present in the network.

4. ANOMALY DETECTION ALGORITHM

The presence of ERE's is an anomalous behavior shared among the problems we target. Because packet loss affects the opening phase of a new TCP connection, RTO timers utilize default values. This introduces well-defined correlations in misconfigured flows at precise time scales dictated by the exponential back-off RTO management algorithm [14]. Thus, if we observe a packet in the three-way handshake that subsequently is lost, then we will observe again the same packet after $3 \cdot 2^k$ seconds ($k=0, 1, 2, \dots$). In principle, if the retransmission sequence were an infinite series, the traffic pattern would produce a power-law ON-OFF behavior known as pseudo self-similarity [5]. However, we observe that, in practice, the sequence of retransmission events is finite and the number of retransmission attempts is limited. In fact:

- TCP/IP module implementations will attempt resending the lost packet a limited number of times. The maximum number of attempts (k_{MAX}) varies. Additionally, k_{MAX} can depend on the state of the TCP connection when the loss occurs (connection opening vs. data exchange). k_{MAX} is typically lower during the handshake stage. For Windows-based hosts in the default configuration, $k_{MAX} = 1$ for the loss of a packet within the handshake phase. In the case of Linux O/S, $k_{MAX} = 4$.
- End-user tolerance to low responsiveness is typically limited to 8-14 seconds [3]. Hence, the TCP module can resend the lost packet only few times before the connection is terminated by the application layer.

Early RTO retransmission patterns (ERE's) are repeated uniformly in all the TCP flows affected by the misconfiguration. This would not be the case if RTO events happened at a later point inside the TCP connection. In fact, RTO timers in each flow would be regulated by the RTT experienced by each connection individually. RTT values typically manifest high dispersion due to the static and dynamic characteristics of the end-to-end connection. RTO retransmissions in the handshake phase are insensitive to such aspects as no RTT measurement is available. Additionally, as default initialization values of the RTO management algorithm are standardized, dependency on a

specific TCP/IP module implementation is not a concern in this phase of the connection.

The steps carried out by our algorithm are as follows. (a.) The MIB packet counter is periodically collected from the link(s) to monitor. This time-series is the input signal to the detection procedure. (b.) The signal is decomposed in the wavelet domain. Wavelets [12] appear to be an ideal instrument because of their built-in scale/time-localization ability. (c.) The algorithm analyzes local minima of the signal energy at resolution levels regulated by the underlying RTO mechanism. (d.) An alarm may be raised to warn administrators and prompt corrective actions. The remainder of this section provides further algorithm details.

Let $\{X_{0,r}\}$ ($0 \leq r \leq 2^M - 1$; $M \in \mathcal{N}$) be the discrete input signal to analyze. The first subscript in $\{X_{0,r}\}$ denotes the aggregation level. The second subscript identifies a specific sample at a given time. Increasing values of the aggregation level correspond to coarser resolutions. ΔT is the constant time interval between two consecutive samples at the finest resolution available. Our algorithm utilizes a Haar-filter based representation of the signal. Two vector series are produced. They are known as the aggregated signals $\{X_{q,r}\}$ (1) and the details $\{d_{q,r}\}$ (2) ($1 \leq q \leq M$):

$$\begin{cases} X_{q,r} = \frac{1}{\sqrt{2}}(X_{q-1,2r} + X_{q-1,2r+1}) & (1) \\ d_{q,r} = \frac{1}{\sqrt{2}}(X_{q-1,2r} - X_{q-1,2r+1}) & (2) \end{cases}$$

Successively, the energy content E_q of the q -th resolution level is computed:

$$E_q = \frac{1}{2^{M-q}} \sum_{r=0}^{2^{M-q}-1} |d_{q,r}|^2 \quad (3)$$

The energy plot is the diagram of $\log_2(E_q)$ as a function of the resolution level q . The detection algorithm uses the energy plot for deriving general aspects of the scaling behavior of the underlying time-series. Asymptotically, the behavior of the energy function is expected to be linear in q for self-similar processes [1] over a broad variety of packet-switched networks:

$$\log_2(E_q) \approx (2H - 1)q + b \quad (4)$$

In (4), H is the Hurst parameter and b is a constant. As $1/2 < H < 1$, the slope of the straight line in (4) is $0 < (2H - 1) < 1$. RTO events alter the linear behavior of the energy function over a precise range of aggregation levels. In our modeling, consecutive RTO events are separated by $3 \cdot 2^k$ seconds ($0 \leq k \leq k_{\text{MAX}}$), being k_{MAX} a finite and generally small value. In the remainder, k_{MAX} is assumed to equal 2.

Locality Property: Let $\Delta T = 3 \cdot 2^u \text{sec}$ ($u \geq 0$) be the signal sampling rate. The energy function of the signal for early RTO retransmissions manifests a local dip over the wavelet aggregation levels $\{u+1, u+2, u+3\}$.

Proof: The signal consists of the initial packet, followed by three subsequent retransmissions. The signal $\{X_{0,r}\}$ can be represented in terms of this binary function: $\delta_0(t) + \delta_{3 \cdot 2^k}(t)$ ($0 \leq k \leq k_{\text{MAX}}$), where $\delta_k(t) = 1$ if $t = k$, $\delta_k(t) = 0$ otherwise. In virtue of (1), the signal at the aggregation level u is:

$$\begin{aligned} X_{u,0} = X_{u,1} = X_{u,3} = X_{u,7} &= 2^{-\frac{u}{2}} & (5) \\ X_{u,2} = X_{u,4} = X_{u,5} = X_{u,6} &= 0 & (6) \end{aligned}$$

The energy content at aggregation levels $\{u+1, u+2, u+3\}$ is:

$$E_{u+1} = \frac{2^{-u}}{4}; E_{u+2} = \frac{2^{-u}}{4}; E_{u+3} = \frac{2^{-u}}{2} \quad (7)$$

Figure 2 shows the plot of the energy function and its shape (local minimum) at such aggregation levels. Figure 2 also contrasts the early RTO-based signal energy function (solid line) with the linear behavior predicted by (4) (dashed line). ■




Figure 2: Energy Function of ERE-Based Signal

In a typical deployment scenario, multiple healthy TCP flows (noise to anomaly detection) will be multiplexed with misconfigured flows (for which the Locality Property holds). The analysis algorithm detects the presence of the misconfigured component embedded in aggregated traffic by studying the energy function shape over an aggregation range inclusive of the interval $[u+1, u+3]$. To locate a dip (local minimum) in the aggregation interval of interest, the energy function is approximated in terms of the least-squares parabola: $y = \beta_0 + \beta_1 x + \beta_2 x^2$. The unknowns $\{\beta_0, \beta_1, \beta_2\}$ are subject to the following conditions:

$$\frac{\partial}{\partial \beta_k} \left(\sum_{i=u}^{u+4} (\log_2(E_i) - \beta_0 - \beta_1 i - \beta_2 i^2)^2 \right) = 0, \quad 0 \leq k \leq 2 \quad (8)$$

Let $\{\tilde{\beta}_0, \tilde{\beta}_1, \tilde{\beta}_2\}$ be the solution to (8). Let V be the vertex of y :

$$V = (V_q, V_{\log_2(E)}) = \left(-\frac{\tilde{\beta}_1}{2\tilde{\beta}_2}; \tilde{\beta}_0 - \frac{\tilde{\beta}_1^2}{4\tilde{\beta}_2} \right) \quad (9)$$

If V satisfies relationships (10) and (11), the detection algorithm marks the time-series as containing an energy dip and, therefore, a sign of anomaly is detected.

$$\begin{cases} \tilde{\beta}_2 > 0 & (10) \\ (u+1) \leq \left(-\frac{\tilde{\beta}_1}{2\tilde{\beta}_2} \right) \leq (u+3) & (11) \end{cases}$$

(10) requires that V is a local minimum. (11) implies that the abscissa of V falls in the energy level range of interest.

Our implementation of the detection algorithm consists of two phases conceptually distinct: a.) wavelet analysis and b.) alarm generation. In the former phase, a time-series of M MIB values is gathered at regular intervals ΔT and wavelet analysis is carried out as described in (8). If relationships (10) and (11) are satisfied, the sample is marked as anomalous. In the alarm generation phase, N time-series are aggregated into a sample $S_{M,N}$. A threshold is used to decide whether to issue an alarm to prompt corrective actions. Let n ($n \leq N$) be the number of time-series in $S_{M,N}$ marked as anomalous and γ be a constant ($0 < \gamma < 1$). An alarm is issued when a sample $S_{M,N}$ is such that $(n/N) \geq \gamma$. The decoupling of the wavelet analysis from the alarm generation algorithm supports enhanced flexibility and improved resiliency to transient conditions, allowing for reduction of spurious alarms.

5. ALGORITHM EVALUATION

We employed two complementary evaluation approaches: simulation and live traces from one of the networks we surveyed.

The detection algorithm consistently utilized the following parameter configuration in both cases: an individual time-series consists of 256 MIB values ($M=256$); MIB sampling time (ΔT) is 1 sec.; a sample $S_{M,N}$ groups 4 time-series ($N=4$); $\gamma = 0.6$. A sliding-window with a width of M seconds is used. After initially filling the window with the first M MIB values, successive window updates refresh historic data by adding 14 new MIB values for each window shift. This configuration allows us to contain detection time below 5 minutes.

5.1 Simulation of Network Misconfigurations

We have designed and implemented a virtual testbed from which to collect data. J-Sim [11] is the discrete-event simulation platform we utilized. These are the steps we followed.

Traffic Model: Initially we extended the simulation platform to support the generation of Web-based traffic. The user behavior is modeled in terms of ON-OFF TCP streams [2]. ON intervals correspond to content transfer from a server to a user. Inactivity periods (OFF intervals) separate consecutive activity periods (ON intervals). ON-OFF times are random values from Pareto distributions. We use the concept of user population to model the aggregation of users that share the same access point. Different user populations model traffic sourcing from different access points. Users' requests inside each user population are distributed over groups of different servers. End-to-end performance metrics (RTT, bandwidth...) vary based on the location of the user, the server and traffic load. This provides for heterogeneity in the aggregated traffic mix. Our testbed incorporates relevant topology and architecture traits of one of the large-scale networks we surveyed.

Preliminary Validation: After extending J-Sim, and before introducing misconfigurations, careful testbed validation was carried out in order to ascertain that healthy aggregated traffic presents statistical properties conforming to expectations [1, 2]. Among the validation tests that we carried out, we measured the degree of self-similarity in terms of the Hurst parameter (H) by using the Abry-Veitch estimator [1]. Different load and simulation conditions were considered. Typically, each simulation gathered 6 hours worth of MIB data. The numerical values that we obtained for H are in the range [0.71:0.83]. Our simulation shows a satisfactory agreement with results in the literature.

Detection Algorithm – Healthy Traffic: The simulation data gathered during the preliminary validation phase was used also to assess the ability of the detection algorithm to correctly identify healthy traffic as such. In fact, no alarms should be raised in this setup for any sample $S_{M,N}$. If an alarm is raised, we incur a false positive. The evaluation criterion we use is to measure the rate of false positives, i.e. the ratio of alarms (number of samples $S_{M,N}$ such that $n/N \geq \gamma$) to the total number of samples. In case no misconfiguration is present, the alarm rate should be low. The second row in Table 1 (part I) shows the numerical results we obtained. The second column indicates the network state, namely no misconfiguration in this case. The third column is the ERE-signal-to-noise ratio expressed as percentage of link throughput due to misconfigured traffic vs. the overall link throughput. A value of 0% signifies that none of the monitored traffic was misconfigured. The fourth column reports the alarm ratio. A value of 0% indicates that no alarms were raised by the detection algorithm. This is in agreement with the fact the network indeed presents no anomalies. Different load conditions were considered

in various simulations. In all cases, the results were consistent with Table 1. Hence, false positives are not a concern.

Detection Algorithm – Misconfigured Traffic: Successively, we evaluated the detection algorithm when anomalies are present. Misconfigurations are initially introduced in the testbed before starting data collection. Misconfiguration modeling closely follows the analysis of section 3. Table 1 (part II) summarizes our findings. Data show that the presence of misconfigurations is strongly correlated with alterations in the shape of the energy function of aggregated traffic at the predicted aggregation levels. The detection algorithm is capable of identifying the presence of an anomalous traffic component in the various misconfiguration cases we target. A high alarm ratio corresponds to effective detection. In the first two cases, the percentage of misconfigured traffic is high enough that all measurements correctly trigger an alarm. On the other hand, the detection algorithm appears to be quite sensitive to traffic anomalies, as the remaining portion of Table 1 shows. In fact, when the relative amount of misconfigured traffic is limited to 0.80% of the total link traffic (fourth case), alarms are correctly reported in more than 53% of the measurements. (The two cases of “Address duplication” differ only in the number of hosts affected by the misconfiguration.)

Table 1: Simulation-Based Algorithm Evaluation

	Network State	Signal-to-Noise Ratio (%)	Alarm Ratio (%)
I	No Misconfigurations	0.00	0.00
II	Dup. Of IP Address Space	100.00	100.00
	Routing Loop	89.45	100.00
	Dup. Of IP Address Space	3.34	69.23
	D-DoS	0.80	53.85

Table 2: Trace-Based Algorithm Evaluation

	Network State	Signal-to-Noise Ratio (%)	Alarm Ratio (%)
I	No Misconfigurations	0.00	9.52
II	Dup. Of IP Address Space	9.23	95.83

5.2 Analysis of Network Traces

Further evaluation of the detection algorithm relied on the use of live traces gathered from the network operated by one of the organizations we surveyed. The data provided to us consists of a set of packet-level traces collected over a link connecting an IDC to a Tier-1 network. Such datasets are routinely collected by field engineers for monitoring and troubleshooting purposes. Trace durations range from about 15 minutes to almost 1 hour. Datasets were collected on different dates and times. One dataset contains traffic affected by a network misconfiguration, namely “Duplication of IP Address Space” severely affecting a significant part of the user population. The remaining traces do not contain anomalies. Packet-level traces were pre-processed in order to produce MIB-equivalent statistics, i.e. the count of packets crossing the monitored link. Such MIB time-series were fed to the detection algorithm. Parameter configuration is identical as in the

previous section. Table 2 summarizes our findings. In case no misconfigurations are present (part I), the overall rate of false positives is approximately 9%. A close analysis of the packet-level traces revealed the presence of traffic generated by custom applications specific to the IDC (mostly for data management). The behavior of such applications may fall outside of the criteria we followed in our modeling. This could explain the discrepancy w.r.t. the simulation results (no false positives). However, also in this “live-scenario”, the incidence of false positives is contained and the algorithm clearly differentiates the healthy traffic case vs. the misconfigured scenario (see part II of Table 2). In fact, when a misconfiguration is present in the collected data, it appears that there is a good agreement between the results obtained in terms of simulation (previous section) and real network traces. The algorithm successfully detects the presence of an anomalous traffic component in over 95% of the measurements.

6. RELATED WORK

An abundance of monitoring infrastructures and misconfiguration detection tools is available on the market ([7] is one instance). A benefit of the methodology we are pursuing vs. other algorithms is its ability to simultaneously capture a broad range of anomalies. This derives from the general mechanisms we exploit. In our analysis, RTO events are a unifying modeling factor shared by a family of misconfigurations. Additionally, several anomaly detection algorithms can be computationally intensive. This restricts their applicability to off-line, post mortem data analysis. Conversely, the algorithm we discuss is relatively simple to implement and integrate in a well-established monitoring infrastructure, i.e. SNMP-MIB.

The work of Huang et al. [8] established the utility of wavelet analysis in inferring network performance parameters, specifically RTT values. However, RTT measurement via analysis of aggregated traffic presents several obstacles such as (a.) high sampling rates due to the typically small RTT values and (b.) high dispersion of RTT values over a broad range. Consequently, the coherence of the multiplexed RTT-based signals to detect is diminished and measurement effectiveness is hindered. The retransmission mechanism we exploit, instead, does not manifest such limitations. Early RTO events happen at connection establishment and default values are used regardless of (a.) the TCP module implementation [14] and (b.) the end-to-end path properties (this is certainly not the case for RTT values). Hence, the embedded signal we use for misconfiguration detection manifests high coherency and allows for effective isolation of anomalies at coarse time scales.

7. CONCLUSIONS

This paper initially classified misconfigurations observed in a group of commercial networks. This provided a concrete snapshot of real-world network problems. Successively, we selected a family of specific misconfigurations from the set of anomalies identified in our survey. Our analysis revealed that over 33% of the problems logged can be modeled in terms of Early RTO Events (ERE’s). Because packet loss occurs in the opening phase of the TCP connection, retransmission timers utilize default values. We utilized the correlation structures induced by ERE’s to isolate the presence of anomalies embedded within the aggregated

traffic mix. Our assessment of this methodology to misconfiguration and anomaly detection is promising. The proposed wavelet-based algorithm effectively identifies the misconfigured traffic component. Therefore, our approach appears to provide a general misconfiguration detection framework. Currently, we are progressing to further evaluation and extensions that can improve the effectiveness of the detection algorithm. To this end, we are investigating the deployment of MIB collectors at frequencies in the range of seconds. Other measurement projects [15] are independently investigating this aspect. Additionally, we are extending the set of misconfigurations that can be modeled by RTO-based events.

8. ACKNOWLEDGMENTS

We would like to express our gratitude to the anonymous referees for their valuable feedback, to Mr. Nomura and all the other colleagues in Fujitsu who provided us with insightful guidance and made this work possible.

9. REFERENCES

- [1] P. Abry, D. Veitch, Wavelet Analysis of Long Range Dependent Traffic, *IEEE Transactions on Information Theory*, Vol. 4, num. 1, 1998.
- [2] P. Barford, M. Crovella, Generating Representative Web Workloads for Network and Server Performance Evaluation, *Proceedings of Performance '98/ACM SIGMETRICS '98*.
- [3] A. Bouch, A. Kuchinsky, N. Bhatti, Quality is in the Eye of the Beholder: Meeting Users’ Requirements for Internet Quality of Service, *SIGCHI conference on Human factors in computing systems*, April 2000.
- [4] C. Fraleigh et al., Packet-Level Traffic Measurements from the Sprint IP Backbone, *IEEE Network*, November 2003.
- [5] L. Guo, M. Crovella, I. Matta, How Does TCP Generate Pseudo-Self-Similarity?, *9th International Symposium in Modeling, Analysis and Simulation of Computer and Telecommunication Systems*, August 2001.
- [6] U. Hengartner, S. Moon, R. Mortier, C. Diot, Detection and analysis of routing loops in packet traces, *2nd ACM SIGCOMM Workshop on Internet Measurement*, 2002.
- [7] *HP OpenView*; available at <http://www.openview.hp.com>.
- [8] P. Huang, A. Feldmann, W. Willinger, A Non-Intrusive, Wavelet-Based Approach to Detecting Network Performance Problems, *ACM IMW'01*, November 2001.
- [9] A. Hussain, J. Heidemann, C. Papadopoulos, A Framework for Classifying Denial of Service Attacks, *SIGCOMM'03*.
- [10] *IETF Working Group: IP Performance Metrics (IPPM)*, available at <http://www.ietf.org/html.charters/ippm-charter.html>.
- [11] *J-SIM*, available at <http://www.j-sim.org>.
- [12] S. Mallat, *A Wavelet Tour of Signal Processing*, 2nd Edition, Academic Press, 2001.
- [13] D. Oppenheimer, A. Ganapathi, D. Patterson, Why do Internet services fail, and what can be done about it? *4th USENIX Symposium on Internet Technologies and Systems*, March 2003.
- [14] V. Paxson, M. Allman, Computing TCP’s Retransmission Timer, *RFC-2988*, IETF, November 2000.
- [15] *SNAPP*, available at <http://snapp.sourceforge.net>