2014 IEEE 28th International Conference on Advanced Information Networking and Applications

# On the Use of Traffic Information to Improve the Coordinated P2P Detection of SLA Violations

Jéferson C. Nobre, Lisandro Z. Granville
Institute of Informatics
Federal University of Rio Grande do Sul - Brazil
Email: {jcnobre, granville}@inf.ufrgs.br

Alexander Clemm, Alberto Gonzalez Prieto
Cisco Systems
San Jose, USA
Email: {alex, albertgo}@cisco.com

***Abstract*—Critical networked services are usually regulated by Service Level Agreements (SLAs). In order to ensure SLAs are being met, it is necessary to monitor Service Level Objectives (SLOs). Active measurement mechanisms are usually chosen to perform this monitoring task, which requires measurement probes to be activated in network devices. However, these probes are expensive in terms of computational resources consumption, thus, active measurement mechanisms usually can cover only a fraction of what could be measured, which can lead to SLA violations being missed. Besides that, highly dynamic networking patterns require the ongoing selection of the candidate network destinations for probing and their respective prioritization, a practice that is not well suited for human administrators because configuring the probes is labor-intensive and error-prone. A possibility to improve the detection of SLA violations is the employment of Peer-to-Peer (P2P) technology in order to steer tasks related to a distributed decision making process for probe activation. In this context, a P2P management overlay can be used to coordinate the probe activation and to share measurement results among the network devices. For a node to rely on measurement data from a peer to determine which probes to configure, it needs to know which of the peers are best correlated with itself, *i.e.*, which nodes have the most significance in terms of being indicative of service level violations that might be observed by the node itself. We propose an autonomic P2P solution to coordinate the placement of active measurement probes in large-scale networks. The edge nodes of the network cooperate via a P2P management overlay to determine what destinations should be monitored. Each edge node determines autonomously what destinations to probe considering local measurements and measurement data from other edge nodes. The measurements considered are traffic information from passive measurement results and past service level measurement results from active measurement results. The proposed solution is evaluated using simulation and the results show its feasibility and interesting features.**

## I. Introduction

Networking infrastructures have evolved in size, complexity, and amount of carried traffic in the last years. Critical networked services provided in these networks require service levels to operate properly. These services levels are usually described in Service Level Agreements (SLAs), established between service provider and customer. Service providers can be financially penalized for SLA violations, since customer applications can suffer from service levels being disrespected. Therefore, human administrators need solutions to ensure that SLAs are not being violated, specially considering the network infrastructure layer. To that end, either active or passive network measurements mechanisms take place.

Passive network measurements are performed through the observation of passing traffic inside network devices, usually in the form of IP flows (*e.g.*, Cisco NetFlow [4]). In passive mechanisms, the network traffic is collected in a non intrusive way because no monitoring traffic is created by the measurement process itself. Active network measurements, on the other hand, are performed through the injection of synthetic traffic into the network (*e.g.*, Cisco Service Level Assurance Protocol [3]). In active mechanisms, measurement probes are distributed along the infrastructure to compute the network performance. Active measurement mechanisms usually offer better accuracy than passive measurements, specially considering service levels. As a result, active is preferred over passive measurement in several scenarios regarding SLA monitoring. However, active measurement is expensive in terms of the resources consumed to process the injected traffic.

The total amount of resources required by active measurement probes on all possible network paths is normally prohibitive. As a result, just a (possibly) small subset of probes is employed, thus covering also a subset of all network flows in a given active monitoring scenario. Choosing which particular probes to deploy in a network is then critical. The traditional practice to distribute these probes into the network consists in $i$) relying on the network administrator to determine a coverage objective, $ii$) collecting traffic information to identify the "hottest" points, $iii$) inferring an initial set of locations to activate measurement probes considering this objective, $iv$) evaluating iteratively the measurement results, and $v$) adapting the locations considering these results and changes in the traffic matrix. However, besides being labor intensive for the human administrator, this practice requires a significant human expertise and once set up (after a lot of effort has been spent), it is rarely changed.

In previous work, we showed that it is feasible to embed Peer-to-Peer (P2P) management software inside network devices to control autonomically the activation of probes using past active measurement results [11]. Besides that, it is also possible to coordinate these activations in order to increase

Jéferson C. Nobre is a PhD student at the Federal University of Rio Grande do Sul, Brazil. He conducted this work during a partial doctoral fellowship at Cisco Systems, funded by CAPES Foundation (Ministry of Education of Brazil).

1550-445X/14 $31.00 © 2014 IEEE
DOI 10.1109/AINA.2014.75

the potential number of probed paths [12]. We showed how P2P coordination of probe placement and mutual sharing of measurement results could be leveraged to significantly improve the likelihood of detection of service level violations. However, we did not make use of one of the main sources of information used by human administrators in probe placement: the use of traffic information to identify traffic-driven SLA violations. Passive measurement mechanisms on the local device can be used to determine which of the potential probe destinations are the most relevant, for example because a large volume of traffic, or traffic that is particularly sensitive to service level violations, is directed their way.

In this paper, we present an autonomic P2P solution to coordinate the activation of active measurement probes using management data from active and passive measurement mechanisms. The integration of both mechanisms is used to enable the introduction of traffic information in probe activation decisions. This is feasible today because of the available programmability support in some network devices (*e.g.*, Cisco onePK [2]). The influence of traffic information is twofold: besides helping in the definition of candidate network destinations for probbing, it is also used to prioritize destinations that are closer to violate the network SLA. The main contributions of this work are: $i$) an increase in the adaptability of probe activation (with respect to previous work), which is desirable considering fast changing network environments (*e.g.*, Data Center networks); $ii$) ensuring the probes are directed at risky and relevant destinations, *i.e.*, not wasted on destinations that are either not relevant (very little traffic directed there from the network device itself) or not in jeopardy of violating SLOs; and $iii$) better scalability features since the amount of configuration parameters required for probe activation is reduced. The proposed solution has been evaluated using PeerSim, an event-based P2P simulator [10]. In our experiments, we deployed our P2P code in different topological settings obtained from the Rocketfuel project [15] and data center designs [7].

The remaining of this paper is organized as follows. In Section II, we review the background on passive and active measurement mechanisms. In Section III, our proposed solution is introduced and its associated concepts are described. The experimental evaluation, encompassing simulation setups, is presented in Section IV. Related work on the use of traffic information to control measurement mechanisms is described in Section V. Finally, conclusions and future work are provided in Section VI.

## II. Measurement Mechanisms

Measurement mechanisms are some of the most important tools deployed by network administrators. Several mechanisms can be used to enable network measurement. In general, these mechanisms are divided into 2 groups: passive measurement mechanisms and active measurement mechanisms. In this Section we first cover some passive measurement mechanisms and their main concepts. After that, most prominent active measurement mechanisms are presented.

### A. Passive measurement mechanisms

In passive measurement, network conditions are said to be checked in a non intrusive way because no monitoring traffic is created by the measurement process itself. Passive measurement data can be used for a variety of purposes. Considering the FCAPS model, there are applications on Fault Management(*e.g.*, abnormal traffic behavior), Configuration Management (*e.g.*, capacity planning), Accounting Management (*e.g.*, ISP billing), Performance Management (*e.g.*, bandwidth monitoring), and Security Management (*e.g.*, flow-based IDS). Passive measurement is realized, for example, inside network devices when they observe the passing traffic *flows*.

Flows can be defined as unidirectional sequences of packets that pass through a network device which are grouped according to some common properties. These properties can consider several packets fields, such as source/destination IP address, source/destination port number, layer 3 protocol type, Type of Service (ToS), and size (aggregated number of bytes). Besides that, other information, such as source/destination Autonomous System (AS) and input/output interfaces can also be used to define flows. For the sake of simplicity, we will present passive measurement mechanisms from Internet Engineering Task Force (IETF) and Cisco Systems.

Cisco NetFlow [4] is a widely deployed protocol used to provide network administrators with access to IP flow information from their data networks. In the context of IETF, the IP Flow Information eXport (IPFIX) Working Group has released several documents describing a protocol, based on the version 9 of NetFlow [5]. Figure 1 shows IPFIX arquitecture (which is based on NetFlow architecture) as an example of passive measurement model. In both protocols, network elements (*e.g.*, routers and switches) gather flow data and export NetFlow/IPFIX records to configured receivers. These receivers are known as collectors (or collecting points).

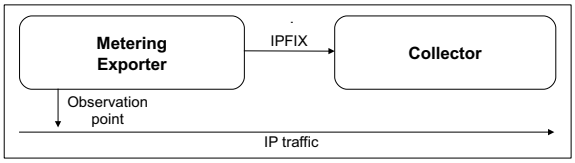


Fig. 1. Passive measurement model

### B. Active measurement mechanisms

Active measurement mechanisms inject synthetic traffic into specific network paths to measure the network performance. These mechanisms can computer the network performance in term, for example, of delay, loss, jitter, and packet/frame loss. Active measurements are performed either one-way or two-way (*i.e.*, round-trip). Two-way measurements, which are common in IP networks, employ time stamps applied at the echo destination to achieve better accuracy.

The generation of synthetic traffic and its computation to provide measurements results is usually performed by an architecture comprised of two hosts with specific roles, a sender and a responder, also collectively known as measurement probes. Figure 2 illustrates this architecture. The exchange of packets between probes is defined by two inter-related protocols: a control protocol, used to initiate and control measurement sessions and to fetch their result, and a test protocol, used to send single measurement packets along the network path under test. The message exchange concerning these protocols is shown in Figure 2. Some of the most prominent active measurement mechanisms are proposed from Cisco and IETF.

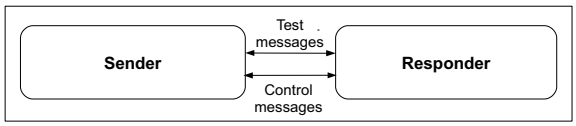


Fig. 2. Condensed active measurement model

Cisco defines the Service Level Assurance (SLA) protocol [3], a broadly known protocol used to measure service level parameters. Cisco SLA protocol measures service levels in layers 2 and 3, as well as applications running on top of layer 3, considering both one-way and two-way metrics. The IETF IP Performance Metrics (IPPM) Working Group proposed open mechanisms that permit the exchange of packets to collect metrics for one-way (One-way Active Measurement Protocol - OWAMP) [14]) and two-way (Two-Way Active Measurement Protocol - TWAMP) [9] packet delay and loss across Internet paths in an interoperable manner.

The employment of active measurement mechanisms is an effective technique for monitoring SLAs and the health of a network as a whole. In order to identify SLA violations using active measurement, it is necessary to have measurement probes activated on problematic end-to-end destinations. However, there is an inherent human and computational cost related to the deployment of these probes and their continuously operation. Thus, improvements in the probe activation decisions can increase the efficiency of SLA violations detection. Since traffic-related issues (*e.g.*, network congestion) have a significant impact on service levels, traffic information can be used to enable better decisions.

## III. The Use of Traffic Information to Improve the Detection of SLA Violations in a P2P Approach

The use of P2P technology in network management (also known as P2P-based Network Management - P2PBNM [8]) can improve the efficiency of SLA violations detection by measurement probes [11]. Embedded P2P management software inside network devices can be used to control probe activation in a decentralized fashion. It is feasible to embed more complex management functions in the devices since they have increased substantially their level of programmability (*e.g.*, Arista EOS Extensibility [1]). In this context, probe activation can be coordinated to increase the maximum number of detected SLA violations since it is possible to share measurement results among devices [12]. Past active measurement results are used to define the P2P management overlay (*i.e.*, which devices are management peers) and to ensure that remote results have local significance. Even though the use of these results can improve the efficiency of SLA violations detection, one source of information is not considered: the traffic matrix.

The traffic matrix is one of the most rewarded information when human administrators plan the deployment of active measurement mechanisms. Traffic information is usually found on network devices in the form of passive measurement results. Since several SLA violations are intrinsically related to congestion and high utilization of network links, it is feasible to have a more efficient SLA violation detection using traffic information. Passive measurement results on the local device can be used to infer the traffic matrix without relying on centralized parties. Thus, it is possible to maintain the desirable decentralized features of a P2P probe activation using local passive measurement results.

In the present work, we devise the utilization of traffic information in order to improve the deployment of active measurement probes by a P2P management overlay. As far as we are aware of, the only studies that incorporate such a overlay to steer the activation of measurement probes were carried out by Nobre *et al.* [11] [12]. In this paper, the use of traffic information by P2P management overlays is twofold: this information is employed to define candidate destinations and to prioritize paths that need to be probed.

In the following subsections, we first discuss the utilization of traffic information from passive measurement mechanisms to select candidate destinations. Then, we explain how coordination strategies for active measurement mechanisms can be improved using the same information to help the formation of a P2P management overlay.

### *A. The Utilization of Traffic Matrix to Select Candidate Destinations*

The traffic matrix is usually employed by human administrators to define candidate destinations for probe activation. Considering the staff expertise, the matrix provides information about which destinations are more "relevant", and are more likely to be communicated with. Hence, a heuristic commonly applied is to prioritize the probbing of these destinations. Thus, probes should be activated manually at top traffic destination by the administrators. In case of changes in traffic matrix, it is necessary to reconfigure manually the probes. This heuristic relies on experienced human administrators for a proper definition of candidate destinations. Besides that, it also consumes significant human resources for continous operation. We propose that the network devices autonomically choose the candidate destinations for probe activation using local passive

measurement results. This information is used to configure the measurement probes, since it is necessary to associate an endpoints for each destination.

Figure 3 depicts schematically the utilization of traffic matrix to select candidate destinations. In this Figure, a network device has traffic information from passive measurement results concerning 4 destinations. The device can activate only 2 probes simultaneously, thus it chooses to probe the ones with more traffic considering its passive measurement results (represented by solid arrows). The number of selected candidates destinations depends of the resource constraints. These constraints represent the number of probes that can be activated in a given time by the network device [11]. The Algorithm 1 describes in more detail the selection of candidate destinations using the traffic matrix. In Algorithm 1, the resource constraints are grouped in a single constant value, $\beta$, which controls the number of selected candidate destinations. The use of a reduced list of candidate destinations can reduce the required resources to bootstrap the probe activation.

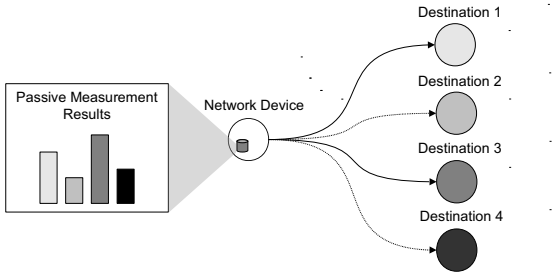


Fig. 3. Selection of candidate destinations by a network device

Passive measurement mechanisms usually offer aggregation capabilities. These capabilities include aggregation of data flows with the same source and destination prefix, source and destination prefix mask, source and destination BGP autonomous system (AS), and input and output interface. One of the most used aggregation scheme is using source and destination prefix. This scheme generates data that allow the examination of the sources and destinations of network traffic passing through the device (described in Section II). Thus, the function $getDestinations$ in Algorithm 1 considers passive measurement results($flowTraffic[]$) in a sliding window ($windowSize$) using source/destination prefix. The sliding window avoids the utilization of "old" results.

Traffic conditions may change over time and the definition of candidates must cope with these changes. Besides that, resource constraints may vary over time. Thus, the candidate destination selection must be performed iteractively. The output of the algorithm 1 is the sorted list of top $M$ candidate destinations ($selectedCandidateDestinations[]$). This list is used as an input to the destination priorization (as described in Subsection III-B)

Algorithm 1 only utilizes local information to prune the list

**Algorithm 1** TrafficCandidateSelection($\beta, windowSize, flowTraffic[]$)

```
candidateDestinations[] ←
getDestinations(flowTraffic[], windowSize, β)
sortDesc(candidateDestinations[])
M ← min(β, sizeOf(candidateDestinations[]))
for i = 1 → M do
    selectedCandidateDestination[i] ←
    candidateDestinations[i]
    i ← i + 1
end for
return selectedCandidateDestination[]
```

of candidate destinations from all possible destinations, defining the most relevant destinations, *i.e.*, most likely targeted by users. The selection of candidate destination is autonomic, in the sense it does not require human intervention, adaptive to changes in network conditions, and independent of the passive measurement technology. In the next Subsection, the list of candidate destinations will be used as the seed for the destination rank and, consequently, to prioritize probe activation.

### *B. Destination Rank and the Prioritization Using Traffic Information*

We consider a scenario of multiple network devices which observe multiple events (end-to-end path measurements), where those devices need to coordinate about the events to observe in a dynamic network. The goal of network devices in our proposed solution is to maximize the number of detected SLA violations. In order to achieve this goal we employ a destination rank, *i.e.*, a sorted list of destinations, to activate measurement probes [11]. This list is dynamically sorted to adapt to changes in network conditions. The use of a destination rank steers a decentralized decision making process since each device holds its own rank. Futhermore, it is possible to use the rank to coordinate the activation of probes [12]. The coordination of probe activation can increase the number of probbed destinations through the sharing of measurements among peer nodes.

The destination rank uses 3 kinds of information to sort the list of destinations: traffic information, past service level measurement results and the time elapsed from the last measurement for a given destination. Each kind of information is included in this rank to a specific purpose. Traffic information is used to define relevant destinations, past service level measurement results capture destinations that are more likely to violate SLAs, and time elapsed from the last measurement is employed to avoid that destinations keep without measurements for a long time. The algorithm 2 describes the decision of process to build the destination rank and the probe activation decision.

Past service level measurement results can be locally collect or received by other network devices. However, it is necessary to assure that the received results have local relevancy.

The assurance of information relevancy uses the concept of *correlated peers* [11]. Two network devices are considered as correlated peers if their measurements for a given destination (or a set of destinations) are correlated. The function $getCorrelatedPeers$ compares the local measurement results with remote ones using Pearson product-moment correlation coefficient. In this context, devices send information about their measurements (collected using a sliding window) for their candidate peers. Eventually, peers also spread their correlated peers in order to permit evaluation of "peers of peers". Correlated peers are also used to enable an autonomic P2P management overlay provisioning.

**Algorithm 2** TrafficCoordinatedDecision($\alpha, \beta, \gamma$, $correlationMin, coordinationMin, windowSize, flows[],$) $selectedCandidateDestination[]$)

```
destinations[] ← selectedCandidateDestination[]
correlatedPeers[]
← getCorrelatedPeers(destinations[], correlationMin)
sendMeasurementSummary(correlatedPeers[])
sendCorrelatedPeers(correlatedPeers[])
shuffle(destinations[])
for t = 1 → sizeOf(destinations[]) do
  rankLast[t] ← getLastLocal(destinations[t])
  rankPast[t] ← getPastLocal(destinations[t])
  rankRemote[t] ← getPastRemote(destinations[t])
  rankTrafLocal[t] ← getTrafLocal(destinations[t])
  t ← t + 1
end for
sortDesc(destinations[], key                    ←
(rankLastLocal[]/ΣrankLastLocal[])              +
(rankPastLocal[]/ΣrankPastLocal[])              +
(rankPastRemote[]/ΣrankPastRemote[])
M ← min(β, α/sizeOf(destinations[],
sizeOf(destinations[]))
for i = 1 → M do
  deployProbe(destinations[i])
  remove(destinations[i])
  i ← i + 1
end for
requestedProbes ← 0
z ← 0
sortDesc(destinations[], key                    ←
(rankTrafLocal[]/ΣrankTrafLocal[])
while requestedProbes < γ and z < sizeOf(destinations[]) do
  candidateCoordinatedPeer ←
  getCorrelatedPeer(destinations[z], coordinationMin)
  sendCoordinationRequest(candidateCoordinatedPeer)
  requestedProbes ← requestedProbes + 1
  z ← z + 1
end while
```

The destination rank is composed by destination scores for each destination. These score are a function of local traffic to the destination ($getLastLocal$), past local measurement results($getLastLocal$ and $getLastLocal$), past peer measurement results ($getPastRemote$). After computing the scores, the destinations on the rank are sorted. The destination rank attempts to build locally a sorted subset that maximizes the fraction of detected SLA violations. After that, according the available local resources, probes are activated on the top destinations. The constraint $\alpha$ is the global upper bound for deployed probes (*i.e.*, concerning nodes that deploy the proposed solution) and $\beta$ is the local upper bound for deployed probes (*i.e.*, in a specific node).

Measurement results for destinations can be shared among multiple devices through coordination of probe activation. In this context, devices exchange messages to "contract" the exchange of measurement results from the chosen devices. The rationale behind the coordination of measurement is that, after activating locally probes at destinations that are likely to violate the SLA, the relevance of destinations according to traffic information is used to decide the ones that should be monitored using measurement result from peers. We call the direct use of remote measurement results as local ones as *virtual probes*. The total number of local virtual probes is controlled by the constraint $\gamma$.

Factors that are taken into consideration by the present solution is that the probability of specific SLA violations and the relevance of these violations. Past service level results can be used to infer the probability of SLA violations. On the other hand, the relevance of probe activation can be improved using traffic information usually from passive measurement results. Despite the fact that passive measurement mechanisms consume substantial results in network decices, their results can improve the probe activation decision as a whole.

## IV. Evaluation

We studied the performance of our proposed solution by defining and implementing simulation experiments. These experiments were implemented in Java using PeerSim [10], an open source event-based simulator of P2P systems. The simulator provides the basic node communication infrastructure as well as transport layer models, which can emulate some characteristics of IP networks (*e.g.*, packet loss and delay). We implemented the probe placement decision algorithms and simple active and passive measurement mechanisms.

The active measurement mechanism shares the basic features found on commercial implementations, such as IPSLA [3]. Thus, it enables the measurement of metrics such as one-way/round-trip delays, jitter, and packet loss in real time. We also implemented a passive measurement mechanism which has features similar to those found on NetFlow [4]. There are specific restrictions on our passive measurement mechanism. First, it considers only collectors (*i.e.*, receivers of flow data) on the own node. Second, we implemented only a "prefix aggregation scheme", thus the flow aggregation is done considering just the prefix of the sources/destinations addresses of network traffic passing through the nodes.

The remaining of the section is organized as follows. First, we describe the experimental setup used to perform the

TABLE I
SUMMARY OF SELECTED TOPOLOGIES USED IN THE EXPERIMENTS

| Topology | Interior nodes | Leaf nodes | Total nodes | Interior/Leaf nodes ratio |
|---|---|---|---|---|
| Rocketfuel WAN | 19 | 21 | 40 | 0.90 |
| "4-post" DC | 20 | 64 | 84 | 0.31 |

TABLE II
EFFICIENCY OF THE SELECTION OF CANDIDATE DESTINATIONS

| Topology | $\beta$ = 1 | $\beta$ =2 | $\beta$ =4 | $\beta$ = leaf nodes |
|---|---|---|---|---|
| Rocketfuel WAN | 5.3% | 10.6% | 21.2% | 100% |
| "4-post" DC | 1.56% | 3.12% | 6.25% | 100% |

simulation experiments. After that, the simulation results are present and discused.

### A. Experimental Setup

We deploy our simulation code using 2 topologies: a WAN topology obtained from the Rocketfuel project [15] (inferred from network environments) and "4-post", a Data Center (DC) topology used and advertised by Facebook [7]. The selected topologies are shown in Figure 4 and some characteristics of these topologies are presented in Table I. We consider that all leaf nodes (depicted as black circles on Figure 4) in these topologies can deploy active measurement probes and passive measurement collectors. The assumption of probes being located on leaf nodes is related to the investigation focus: detection of *end-to-end* SLA violations. These assumptions also holds considering the common practices on field deployments.

For simplicity, we used a network-wide SLO for detecting SLA violations. However, it is possible to operate with multiple SLAs as well as considering SLOs in a per node basis. All network modifications are injected in one direction for a given link and are defined in terms of simulation cycles. In the experiments, we used series of 10 simulation experiments; the observed variance in the experiments was low.

### B. Results

The focus of the simulation experiments is to evaluate the utilization of traffic information considering the detection rate of SLA violations. According to the proposed solution (described in Section III), traffic information can help the selection of candidate destinations and the prioritization of relevant destinations. We explicitly do not focus on the accuracy characteristics of the measurement mechanisms themselves, since we believe our approach can be used with different measurement mechanisms.

Initially, we aim at determining the efficiency of the selection of candidate destinations considering the topologies presented in IV-A. In order to accomplish that, first we show the percentage of selected destination (*i.e.*, $\beta$ - the number of local activated probes) against the complete list of destinations on Table II for the Rocketfuel WAN topology and the "4-post" DC topology. The values on this Table were chosen to present the worst case scenario, when the space of all possible destination has to be explored using just a few probes. We also include the maximum number of leaf nodes of each topology to highlight that in this case the selection efficiency is not interesting.

The utilization of traffic information can improve the efficiency of the selection of candidate destinations considering the selected topologies. This information can be used to select "top" traffic destinations which are usually the more relevant ones. Besides that, if there are traffic-related SLA violations (*e.g.*, due to congestion), the adaptation of the probe activation mechanism can be faster, specially considering larger infrastructures. For example, regarding the selected topologies, this is more important on the second one since it has a larger destination space to cover. In fact, excluding the situation where the available resource are sufficient to probe all possible destinations ($\beta$ = leaf nodes) and obviously there is no impact on using traffic information, the use of this information leads to a smarter employment of resources.

We performed experiments to determinate the adaptation features of the utilization of traffic information to prioritize probe activation. In order to accomplish that, we collected the total number of SLA violations detected by nodes regarding a specific network environment setup (for the mentioned topologies). In this setup, we increased the one-way delay on 4 access links for 20 cycles, and then we changed for other 4 links for the same amount of cycles. This increase makes the end-to-end paths that traverse the changed links to appear as SLA violating for the simulated active measurement mechanism. We chose the number of cycles in which the experimental setup is changed in order to permit that the proposed approaches go through their permanent response. The traffic is equally divided among all the endpoints and it is constant. We show on Figure 5 results for experiments performed on the Rocketfuel WAN topology and on Figure 6 results for experiments performed on the "4-post" DC topology.

The curves depicted on Figure 5 and Figure 6 represent the mean number of detected SLA violations as a function of simulation cycles. Besides that we represent the relevance of the violations using the product of detected violations and the traffic directed to the violating endpoints ("relevance" on Figures). The number of locally deployed probes ($\beta$) is 3 and the maximum number of coordinated measurements ($\gamma$), *i.e.*, virtual probes, is set to 1 per node. Both Figure show results obtained without the use of traffic information (Figures 5(a) and 6(a)) and with the use of traffic information (Figures 5(b) and 6(b)). Besides that, we also present the maximum number of SLA violations that can be detected by local probes ("beta max" on Figures) , *i.e.*, every local probe detects a SLA violation, and the maximum number of SLA violations that can be detected by local and virtual probes ("max" on Figures), Both "beta max" and "max" are defined considering the number of probes that can be activated by each node and the total number of nodes that can activate probes (leaf nodes).

As can be seen in both Figures 5 and Figure 6, the utilization of traffic information can improve the relevance of detected

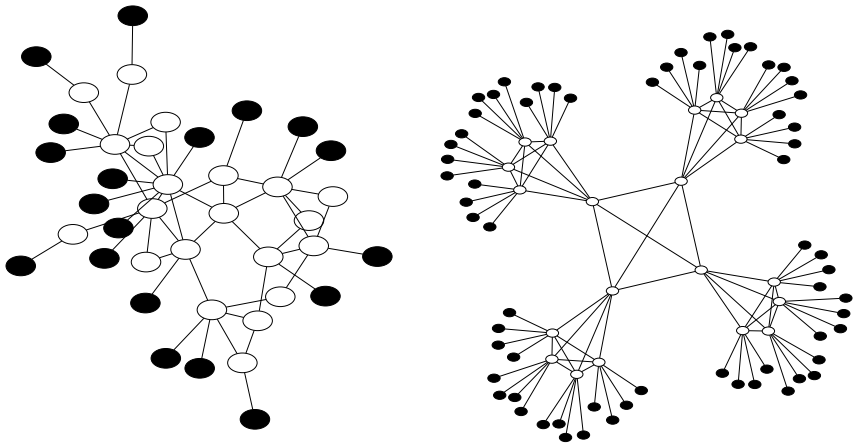

(a) Rocketfuel WAN topology (b) "4-post" data center topology

Fig. 4. Selected topologies used for simulation experiments.

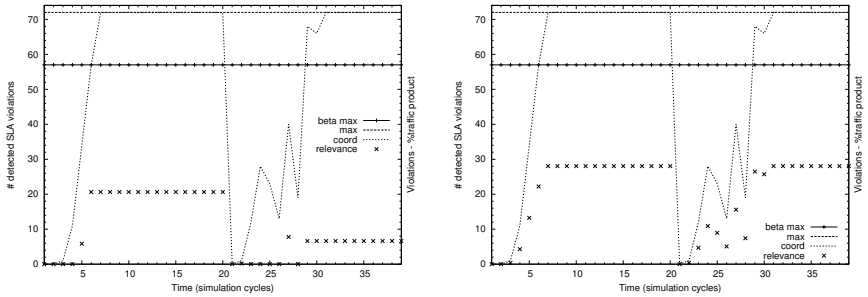


(a) Without traffic information (b) With traffic information

Fig. 5. Results for Rocketfuel WAN topology.

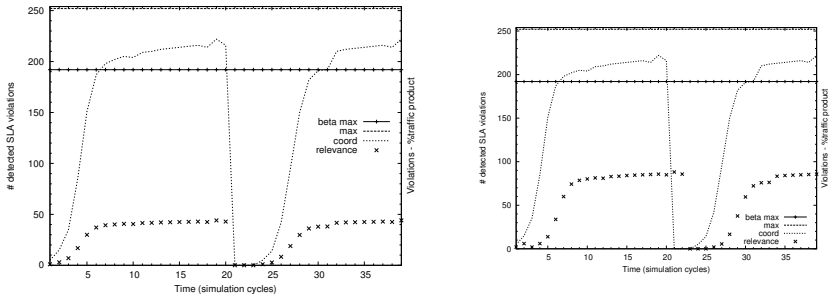


(a) Without traffic information (b) With traffic information

Fig. 6. Results for "4-post" DC topology.

SLA violations without decreasing the total number of detections (for the aforementioned topologies and setups). Clearly, even the utilization of traffic information to steer the choice of just 1 coordinated measurement is positive. Besides that, we approached almost the maximum number of detected SLA violation for the resource constraints. The experiments also show that the proposed approaches behave as we expected, without stability and convergence problems. Therefore, if there is a sufficient number of probes (considering the number of violations), the coordinated algorithm will find them and converge.

It is worth mentioning that there is a trade-off between the utilization of the traffic information and better results. Since this information usually come from passive measurement mechanisms, the resources required to deploy these mechanisms should be also taken into account. In most infrastructures passive measurement results are used to perform several management tasks, such as traffic engineering, thus it is feasible to consider that traffic information is ordinarily available. In any case, the human administrator could also

use sampled results from the passive measurement mechanism (*e.g*, Sampled Cisco NetFlow), a methodology that decreases the resource consumption.

## V. Related Work

The use of traffic information to tacke the problem of measurement probe assignment was investigated in the context of some research initiatives over the past years. Some of these initiatives are discussed as follows.

Sekar *et al.* [13] proposed CSAMP, a centralized optimization engine for system-wide flow monitoring. The main features of CSAMP are the use of traffic information to steer flow sampling and hash-based packet selection through a centralized engine for the distribution of measurement responsibilities across routers. The authors claim that CSAMP can provide greater monitoring coverage and an improved use of router resources. However, as CSAMP relies on a centralized party there are concerns about reliability as well as the cost and delay associated with the dissemination of routing manifests. Furthermore, the approach requires modifications in the measurement mechanisms.

Pietro *et al.* [6] proposed DECON, a decentralized coordination system aimed at assigning passive monitoring probes. DECON uses traffic information from probes seeing a particular ow to decide which one shoud do the actual monitoring. After that, messages are sent back to probes communicating the decision. Authors claim that DECON scales up to large numbers of flows without requiring network topology information nd packet marking. However, DECON operates using a detached P2P overlay, thus it is necessary to add up the ownership cost of additional hardware due to the detached overlay.

## VI. Final Remarks and Future Work

Critical networked services established between service provider and customer are expected to operate respecting Service Level Agreements (SLAs). An interesting possibility to monitor these SLAs is using active measurement mechanisms. However, these mechanisms are expensive in terms of network devices resource consumption. Hence, only a subset of end-to-end destinations is actually measured in most network infrastructures. This can lead to SLA violations being missed, which invariably affects the performance of several applications. The current best practice, the observation of just a subset of network paths driven by the expertise of the human administrators, is error prone and does not scale well. In addition, the SLA violations are more relevant when they occur in heavily utilized links since this impairs the service levels of more users.

In this paper we propose the use of traffic information from passive measurement results in an autonomous and dynamic way to help the definition of the candidate destinations that need probe activation and to prioritize these destinations. Our solution is based on the utilization of a Peer-to-Peer (P2P) management overlay to enable the sharing of managenement information among network devices. The contribution of our proposal is to provide a solution for decentralized control of probe activation that is adaptive to changes in network conditions, does not require human intervention, and is independent of underlying active and passive measurement mechanisms. Furthermore, we have presented an evaluation of our proposed solution using simulation. The results show that traffic information can be used to improve probe activation decisions.

Although the proposed solution shows good results in simulation experiments performed until the present moment, we intend to investigate different conditions (*e.g.*, using data from real network traces). We are also working on to develop and evaluate a validation phase for probe activation decision. In this phase, simpler measurement probes would be activated just to enable agreements among peers about coordination of destinations to be actually measured. Besides that, information about correlated/coordinated peers can have other uses. For example, it can be used to infer the underlying network topology.

## References


[1] "Arista EOS Extensibility." [Online]. Available: http://eos.aristanetworks.com/wiki/index.php/EOS_Extensibility

[2] "Cisco One Platform Kit." [Online]. Available: http://developer.cisco.com/web/onepk/home

[3] M. S. Chiba, A. Clemm, S. Medley, J. Salowey, S. Thombare, and Y. E., "Cisco service-level assurance protocol," RFC 6812 (Informational), Internet Engineering Task Force, January 2013.

[4] B. Claise, "Cisco systems netflow services export version 9," RFC 3954 (Informational), Internet Engineering Task Force, October 2004.

[5] ——, "Specification of the ip flow information export (ipfix) protocol for the exchange of ip traffic flow information," RFC 5101 (Standard), Internet Engineering Task Force, January 2008.

[6] A. di Pietro, F. Huici, D. Costantini, and S. Niccolini, "Decon: Decentralized coordination for large-scale flow monitoring," in *Proceedings of the IEEE Conference on Computer Communications (INFOCOM) Workshops*, march 2010.

[7] N. Farrington and A. Andreyev, "Facebook's data center network architecture," Research Publications at Facebook, Facebook, Inc., 2013.

[8] L. Z. Granville, D. M. da Rosa, A. Panisson, C. Melchiors, M. J. B. Almeida, and L. M. R. Tarouco, "Managing computer networks using peer-to-peer technologies," *Communications Magazine, IEEE*, vol. 43, no. 10, pp. 62–68, 2005.

[9] K. Hedayat, R. Krzanowski, A. Morton, K. Yum, and J. Babiarz, "A two-way active measurement protocol (twamp)," RFC 5357 (Proposed Standard), Internet Engineering Task Force, October 2008.

[10] A. Montresor and M. Jelasity, "PeerSim: A scalable P2P simulator," in *Proceedings of the 9th International Conference on Peer-to-Peer (P2P)*, 2009.

[11] J. C. Nobre, L. Z. Granville, A. Clemm, and A. G. Prieto, "Decentralized Detection of SLA Violations using P2P Technology," in *Proceedings of the 8th International Conference on Network and Service Management (CNSM 2012)*, October 2012.

[12] ——, "Coordination in P2P Management Overlays to Improve Decentralized Detection of SLA Violations," in *Proceedings of the IEEE International Conference on Communications (ICC 2013)*, (June 2013.

[13] V. Sekar, M. Reiter, W. Willinger, H. Zhang, R. Kompella, and D. G. Andersen, "Csamp: a system for network-wide flow monitoring," in *Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, april 2008.

[14] S. Shalunov, B. Teitelbaum, A. Karp, J. Boote, and M. Zekauskas, "A one-way active measurement protocol (owamp)," RFC 4656 (Proposed Standard), Internet Engineering Task Force, September 2006.

[15] N. Spring, R. Mahajan, and D. Wetherall, "Measuring isp topologies with rocketfuel," *ACM SIGCOMM Computer Communication Review*, vol. 32, no. 4, pp. 133–145, 2002.