Coordination in P2P management overlays to improve Decentralized 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 enable significant revenue for network operators and, in turn, are regulated by Service Level Agreements (SLAs). In order to ensure SLAs are being met, service levels need to be monitored. One technique for this involves active measurement mechanisms which employ measurement probes along the network to inject synthetic traffic and compute the network performance. However, these mechanisms are expensive in terms of 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, the definition of this fraction is a practice done by human administrators, which does not scale well and does not adapt to highly dynamic networking patterns. The contribution of the present work is the proposal of a solution to increase the potential number of detected SLA violations in which network devices autonomously and dynamically share service level measurement results. The sharing of these measurement results is based on the utilization of a Peer-to-Peer (P2P) management overlay built using past service level measurement results and a coordination strategy characterized by a high degree of decentralized decision making. The solution is evaluated using simulation and the results show its feasibility and interesting features.**

I. Introduction

Computer networks have evolved in size and complexity in the last years. Besides that, the service level requirements of critical networked services provided in these networks have become a critical concern for network administrators. In this scenario, these services are expected to operate respecting associated Service Level Agreements (SLAs), established between service provider and customer. In this context, solutions to ensure that SLAs are not being violated are crucial. However, predominant solutions rely on a heavy utilization of computational and human resources, which sometimes cannot be met.

Active measurement mechanisms are the prime choice for SLA monitoring. In these mechanisms, measurement probes are distributed along the network to inject synthetic traffic and compute the network performance. However, active measurement is expensive in terms of the resources consumed to process the injected traffic and deliver the SLA metrics. In addition, the total amount of resources required for probe deployment is strongly related to the number of possible network paths. Besides that, if the critical flows are too short in time and dynamic in terms of traversing network paths (like in modern cloud environments), fast reactions are necessary to reconfigure the probes. In this context, active measurement is usually activated only on a subset of all available probes by the network operator, *i.e.*, the set of all network flows is never covered entirely. Furthermore, this subset does not keep up with dynamic network environments.

The current best practice to distribute the available measurement probes along the network consists in relying on the network administrator to collect traffic information and infer which are the best locations to deploy measurement probes. However, besides being too difficult and labor intensive for the human administrator, this practice is inefficient considering fast changing network environments. Some works [6] propose to embed management software inside network devices to control probe distribution using a P2P management overlay. This is feasible since these devices have increased substantially their level of programmability (*e.g.*, Cisco onePK). However, the number of detected SLA violations is still bounded by the number of deployed probes. Thus, if the number of SLA violation is greater than the number of available probes, only a fraction of the violations would be observed, which invariably leads to the damaging consequence of undetected SLA violations.

In this paper, we present an autonomic Peer-to-Peer (P2P) solution to coordinate the deployment of active measurement probes which increases the number of detected network SLA violations over the available measurement probes. The solution introduces features inspired by multi-agent coordination, a concept borrowed from Multi-Agent Systems (MAS). These features steer a decentralized decision making process to determine which measurement probes must be activated/deactivated to cope with the network dynamics. Our solution is adaptive to changes in network conditions, independent of the underlying active measurement technology, and requires no human intervention. The proposed solution has been evaluated using an event-based P2P simulator. In our experiments, we deployed our peer code in different topological settings obtained from the Rocketfuel project [9]. Results show that our solution performs better than a non-coordinated probe placement, decreasing the number of potentially missed SLA

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).

978-1-4673-3122-7/13/$31.00 ©2013 IEEE

violations.

The remaining of this paper is organized as follows. In Section II we review the background on active measurement mechanisms and coordination strategies for distributed network management tasks. 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. Finally, conclusions and future work are provided in Section V.

II. Background

In this Section we first cover some of the most prominent active measurement mechanisms and their main concepts. After that, some strategies commonly used to coordinate distributed systems are presented.

A. Active measurement mechanisms

Active measurement mechanisms inject synthetic traffic into specific network paths to measure the network performance which can lead to either one-way or two-way (*i.e.*, round-trip) measurements. The generation of synthetic traffic and its computation to provide measurements is usually performed by an architecture comprised of two hosts with specific roles, a sender and a responder, also collectively known as measurement probes. Active measurement mechanisms have good characteristics in terms of accuracy and privacy [7]. However, these mechanisms are expensive specially in terms of computational resources consumed on network devices. The amount of consumed resources explodes with the size and complexity of current network infrastructures which avoids monitoring all network flows. There are several protocols used to enable active measurement, however, for the sake of simplicity, we will focus active in the proposals from IETF and leading networking equipment vendors.

The Internet Engineering Task Force (IETF) IP Performance Metrics (IPPM) Working Group has proposed open mechanisms that permit the exchange of packets to collect metrics for one-way (One-way Active Measurement Protocol - OWAMP) [7]) and two-way (Two-Way Active Measurement Protocol - TWAMP) [3] packet delay and loss across Internet paths in an interoperable manner. These mechanisms consist of 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 Internet path under test.

From the industry side, it is possible to highlight solutions from some network equipment vendors. Juniper Networks presents Real-time Performance Monitoring (RPM) [1] to enable the configuration of active probes in order to track and monitor traffic across the network, and investigate network problems. The main use for RPM is performance monitoring on layers 3 and 4 and it can also generate SNMP traps on SLA violations. Cisco Systems defines the Service Level Assurance (SLA) protocol[2], a widely deployed 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, both considering one-way and two-way metrics. Both protocols support a variety of network statistics including delay, loss, jitter, and packet/frame loss. Besides that, these protocols are known to consume significant amounts of resources what leads to the use of dedicated routers (also known as "shadow routers") to handle active measurement mechanisms and save resources for primary network functions (*e.g.*, routing and switching).

B. Coordination strategies for distributed network management tasks

Coordination strategies can be used as a basic building block for distributed network management tasks. However, the effective implementation of these strategies poses significant theoretical and practical challenges. Therefore, for coordination strategies to be successful, numerous issues must be addressed, including the definition and management of shared information among a group of management nodes to facilitate the coordination of these nodes. Information necessary for coordination may be shared in a variety of ways. However, for the sake of simplicity, we will focus only in a few coordination strategy proposals.

Yalagandula *et al.* [11] proposed S^3, a Scalable Sensing Service (thus, the "S^3" acronym) aimed at SLA management, resource provisioning, and fault management. This service uses *Scalable Distributed Information Management System* (SDIMS) [10] as "sensing backplane" to coordinate implicitly monitoring tasks performed by different *sensor pods* (a collection of lightweight measurement and monitoring sensors that collect information at a machine) through a publish-subscribe scheme. S^3 is developed by Hewlett-Packard Laboratories and is being used in a recently funded DARPA Internet Control Plane project called CHART (Control for High Throughput Adaptive and Resilient Transport).

Nobre *et al.* [5] proposed the utilization of multi-agent truth maintenance features concepts to introduce consistency data maintenance and coordination features in P2P-based autonomic network management systems. The authors apply belief propagation among nodes to coordinate explicitly management tasks. These beliefs are spread over the P2P management overlay using a proliferation model, which is drawn from biology-inspired distributed computing models. The authors describe 2 case studies regarding network management policies [5] and collaborative fault management of links in access networks [4].

III. Coordinating nodes to deploy measurement probes

Active measurement mechanisms are an effective technique for monitoring Service Level Objectives (SLOs) and the health of a network as a whole. These mechanisms can be

[1]Real-Time Performance Monitoring on Juniper Networks Devices - http://www.juniper.net/us/en/local/pdf/app-notes/3500145-en.pdf

[2]Cisco IOS IP Service Level Agreements - http://www.cisco.com/go/ipsla

employed in different contexts, such as pre-deployment validation or measurement of in-band live network performance characteristics. However, these mechanisms are expensive in terms resource consumption, both on network devices and bandwidth. The reason for this is the inherent cost related to the generation and forwarding of synthetic test packets and the further analysis of those packets and their responses.

There are different approaches to improve the efficiency of the deployment of measurement probes. Some works [6] propose to embed P2P management software inside network devices to control probe distribution with positive results concerning the rate of detected SLA violations. The introduction for P2P software inside network devices is feasible since these devices have increased substantially their level of programmability (*e.g.*, Juniper Junos Script Automation). However, the number of detected SLA violations is still bounded by the number of deployed probes. Thus, if the number of SLA violations is greater than the number of available probes, only a fraction of these violations would be observed.

In the present work, we devise the use of coordination strategies among peers to steer probe placement decision, *i.e.*, to decide which paths must be probed observing resource constraints. Peers exchange management messages to coordinate the paths they deploy measurement probes. Thus, our conceptual solution is based on P2P technology and is characterized by a high degree of decentralized decision making across network nodes. Distributed decision making has an advantage, in this context, over centralized decision making in the sense that a peer is not required to access information from all the other peers (to make a decision).

Network devices in our solution coordinate to define which end-to-end paths must be measured locally and which ones can be remotely probed. First, our solution attempts to build locally a subset that maximizes the fraction of detected SLA violations. This subset is built using information from past service level measurement results performed locally (*i.e.*, produced by the own device) or received from other network devices. Second, network devices rank the paths in this subset what leads to a distributed path rank. After that, this rank and the resource constraints are used to select a final subset of end-to-end paths that must be probed locally as well as those whose results are expected to be received from other devices. Finally, devices exchange messages to "contract" the selected paths from the chosen devices.

In the following subsections, we first discuss the concept of correlation peers and consequently the P2P management overlay formation. Then, we explain the coordination strategies proposed to improve the capabilities of the active measurement mechanisms deployed by the network devices

A. Correlated peers

P2P technology has several distinctive characteristics that make it interesting for network management [2]. In P2P-Based Network Management (P2PBNM), most of the system states and tasks are directly and dynamically allocated among the peers. The resulting decentralization also pushes the peers' local autonomy, increasing the use of local data and logic to make management decisions [1]. Specially considering in-network (*i.e.*, embedded) P2PBNM system, the decentralization of these systems makes their growth more "organic" since the addition of new devices can also follow the introduction of new management peers. However, the overlay provisioning must be transparent in order to ease this growth.

In the present proposal, the P2P management overlay provisioning uses the concept of *correlated peers* [6]. Two devices are considered as correlated peers if their measurements for a given destination (or a set of destinations) are correlated. In order to bootstrap overlay formation, each peer uses their known endpoints neighbors as the initial seed to get candidate peers, *i.e.*, peers that may be evaluated for correlation purposes. After that, devices send information about their measurements (collected using a sliding window) for their candidate peers. Each device compares this information with their own measurements and candidate peers are then ranked using the obtained correlation scores. Then, the candidate peers with top scores become correlated peers. Eventually, peers also spread their correlated peers in order to permit evaluation of "peers of peers".

There are different ways to compare the locally produced measurement results and those received by other devices in order to define correlated peers. Among the different choices, the authors initially chose to perform this comparison using a *Student's t-test* since they argued this test can be done using only summary information, therefore the information that needs to be exchanged by peers is substantially decreased [6]. However, other kinds of tests can be used to compare results, such as those using the *Pearson product-moment correlation coefficient* and *Spearman's rank correlation coefficient*.

B. Coordination strategy and measurement probe placement

We now present how the management peers coordinate and measurement probe placement algorithms take place in our solution. Concerning the utilization of active measurement mechanisms to detect SLA violation, there is an inherent trade-off between attempting to maximize SLA coverage over end-to-end paths and minimizing the resource consumption. However, even considering maximum coverage (*i.e.*, increasing the number of deployed probes without considering resource consumption), the number of measurements that a device can perform is bounded by the available resources. Our proposed solution to increase the number of SLA violation tries to capture one of the behaviors commonly employed by network administrators, the coordinated sharing of measurement results.

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. Hence some mechanism is needed to dynamically adapt the coordination strategies to network conditions. We use an approach which is opposite to hard coordination. In fact, we employ a computationally simple algorithm that contracts measurements from other devices through message exchange in a P2P management overlay. Besides that, the resource constraints are defined according to the maximum number of probes expected to be deployed in a given time [6]. This number is used as an abstraction for the resources available to deploy active measurement mechanisms. The constraint α is the global upper bound for deployed probes (*i.e.*, concerning nodes that deploy the proposed solution) and β is the local upper bound for deployed probes (*i.e.*, in a specific node).

The coordination strategy to deploy active measurement probes uses shared information which take the form of common objectives. These objectives are the choice of which end-to-end paths must be probed. In order to accomplish this coordination we have developed a simple protocol. This protocol is inspired in the same kind of mechanism found in coordinated decision making protocols from multi-agent research area. For example, the *Contract Net Protocol* (CNP) [8] is a multi-agent task-sharing protocol, consisting of a collection of agents, which can act like a manager or a contractor (*e.g.*, for different tasks). Agent, then, exchange bids and contract are made between managers and contractors. We

Our coordinated strategy enhances the concept of correlated peers in order to choose which peers are interesting to share measurement with. We will use the terminology commonly used by multi-agent coordination to describe our strategy. The goal of network devices in our proposed solution is to maximize the number of detected SLA violations. First, a network device must recognize that it has a problem to achieve this goal and it wants help with. A single device can either realizes it cannot achieve the goal in isolation (*i.e.*, there are more SLA violations than locally available probes), or realizes it would prefer not to achieve the goal in isolation (typically to save resources for main network functions such as switching and routing). Second, the device decides which peer is the best candidate to share measurements. The choice is made considering the top correlated peer (*i.e.*, peer with the highest correlation score). Since the P2P management overlay is built using information from past service level measurement results (previously explained), therefore, indirectly, the probability of a peer to be chosen is related with the same results. In order to start the coordination, a coordination request is sent by the device to the chosen correlated peer. The peers that receive the coordination request can return the request with either a positive coordination response or a negative coordination response.

The correlated peers with coordination measurements need to exchange information, more specifically, the summary of measurement results. However, the network conditions may change over time and the coordination strategy must cope with these changes. For coordination strategy to be effective, the correlated peers must be able to respond to unanticipated situations or changes in the network environment that are sensed as a coordinated task is carried out. These changes can make a peer to stop the measurement for a given end-to-end path which is coordinated with other peer(s). Surely, the peer must inform about this stop to the peers which share the related measurement.

IV. Evaluation

We studied the performance of our proposed solution by defining and implementing simulation experiments. These experiments were implemented in Java using PeerSim[3], 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 as well as a simple active measurement mechanism. The focus of the simulation experiments is to evaluate the coordination strategy considering the detection rate of SLA violations as well as the properties of the resources constraints per se. We explicitly do not focus on the accuracy characteristics of the measurement mechanism itself, since we believe our approach can be used with different active measurement mechanisms.

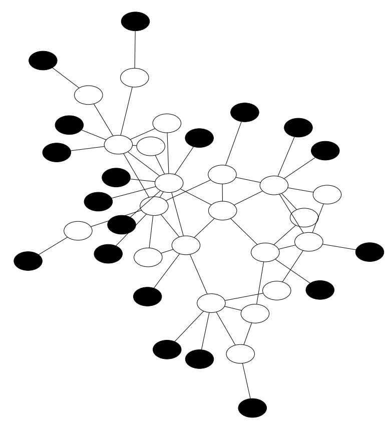

Fig. 1. An example of RocketFuel topology

We deploy our simulation code using topologies inferred from real network environments. These topologies are obtained from the Rocketfuel project [9]. One example of this topologies is shown on Figure 1. We consider that all leaf nodes (depicted as black circles on Figure 1) in these topologies can deploy active measurement probes as well as these probes always reply to measurement requests. 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. All end-to-end paths are considered as candidates to be probed in each simulation cycle.

We show on Figure 2 the potential number of missed SLA violations as a function of the number of locally deployed probes (β) regarding a specific network environment setup. In this setup, we increased the one-way delay in the example topology using 3 different percentages on access links: 10%, 25%, and 50% (considering the local device view). This increase makes the end-to-end paths that traverse the changed

[3]PeerSim: A Peer-to-Peer Simulator - http://peersim.sourceforge.net

links to appear as SLA violators for the simulated active measurement mechanism. The maximum number of coordinated measurements is set to 1 per node. The curves depicted on Figure 2 represent the potentially missed SLA violations considering the changed access links and the utilization of coordination strategy: with the use of coordination strategy ("coord - 0.1" for 10%, "coord - 0.25" for 25%, "coord - 0.5" for 50%) and without the use of coordination strategy ("no coord - 0.1" for 10%, "no coord - 0.25" for 25%, "no coord - 0.5" for 50%).

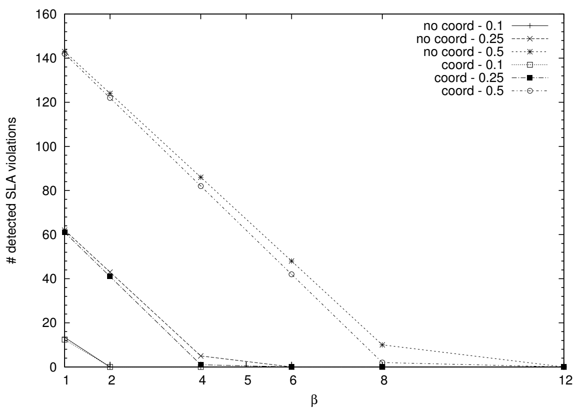

Fig. 2. Number of potentially missed SLA Violations

As can be seen in Figure 2, the utilization of a coordination strategy decreases the number of missed SLA violations in the experimental setup (less is better on Figure 2). Clearly, even the utilization of only 1 coordinated measurement improves the SLA detection performance. Since the coordination protocol (as described in Section III) is extremely simple, it is safe to consider that the trade-off between the coordination overhead (*i.e.*, messages exchanged between correlated peers) and the performance improvement is positive.

V. Final Remarks and Future Work

Critical networked services established between service provider and customer are expected to operate respecting 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. In addition, if the number of SLA violation in a given time is higher than the number of available measurement probes (common place in large and complex network infrastructures), certainly some SLA violations will be missed. The current best practice, the observation of just a subset of network paths driven by human administrators' expertise, is error prone and does not scale well.

In this paper we propose a solution to increase the potential number of SLA violations detected by an active measurement mechanism in a network infrastructure. Our proposed solution aims at the coordination of service level measurement in a per network device basis. Our solution is based on the utilization of a Peer-to-Peer (P2P) management overlay to enable the sharing of measurements among network devices. The coordination strategy uses the concept of correlated peers (*i.e.*, devices that have correlated measurement results) in order to define which peers are good choices to coordinate with. Furthermore, we have presented an evaluation of our proposed solution using simulation. The results show that the coordinated sharing of service level measurements can improve the potential number of detected SLA violations.

Although the proposed solution shows good results in simulation experiments performed until the present moment, it is necessary to evaluate different kind of topologies (*e.g.*, data center networks) and conditions (*e.g.*, using data from real network traces). We also intend to investigate how coordination features can enable composite measurement tasks, *i.e.*, measurements can be break down into sub-tasks (if possible), and, thus lead to the cooperation of peers among the P2P management overlay to further increase the number of detected SLA violations. Furthermore, our ongoing work is aimed at developing a prototype to run on Software-Defined Networking (SDN) equipment which provide an increasing level of processing power and programmability.

References

[1] P. A. P. R. Duarte, J. C. Nobre, L. Z. Granville, and L. M. R. Tarouco, "A P2P-Based Self-Healing Service for Network Maintenance," in *Proceedings of the 12th IFIP/IEEE International Symposium on Integrated Network Management (IM 2011)*, may 2011.

[2] 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.

[3] 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.

[4] J. C. Nobre and L. Z. Granville, "Consistency of States of Management Data in P2P-Based Autonomic Network Management," in *20th Ifip/Ieee International Workshop on Distributed Systems: Operations and Management (DSOM 2009), Venice, Italy. October 2009.*, 2009, pp. 99–110.

[5] ——, "Consistency of Policy States in Decentralized Autonomic Network Management," in *Proceedings of the 12th IEEE/IFIP Network Operations and Management Symposium (NOMS 2010), Osaka, Japan. April 2010.*, 2010.

[6] 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), Las Vegas, USA. October 2012.*, 2012.

[7] 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.

[8] R. Smith, "The contract net protocol: High-level communication and control in a distributed problem solver," *Computers, IEEE Transactions on*, vol. 100, no. 12, pp. 1104–1113, 1980.

[9] 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.

[10] P. Yalagandula and M. Dahlin, "A scalable distributed information management system," *ACM SIGCOMM Computer Communication Review*, vol. 34, no. 4, pp. 379–390, 2004.

[11] P. Yalagandula, P. Sharma, S. Banerjee, S. Basu, and S.-J. Lee, "S3: a scalable sensing service for monitoring large networked systems," in *Proceedings of the 2006 SIGCOMM workshop on Internet network management*, ser. INM '06. New York, NY, USA: ACM, 2006, pp. 71–76.