

Global Vs. Per-Domain Monitoring of Multi-Domain Networks

Emna Salhi, Samer Lahoud, Bernard Cousin
IRISA / University of Rennes 1, France
{emna.salhi, samer.lahoud, bernard.cousin}@irisa.fr

Abstract—Multi-domain monitoring aims at guaranteeing QoS for services crossing several domains. It is often desirable to perform global monitoring to guarantee end-to-end QoS for services across domains and to reduce the monitoring cost. However, global monitoring might be infeasible due to confidentiality constraints. The alternative solution is to perform per-domain monitoring.

In this work, we propose to evaluate global and per-domain monitoring techniques. For this end, we study the properties of multi-domain networks and the requirements of multi-domain monitoring. We formulate the problem as an Integer Linear Program (ILP). We show that it is a Nondeterministic Polynomial Time Hard (NP-Hard) problem, and therefore, we devise a heuristic that meets multi-domain properties. We show that confidentiality is far from being the only constraint to global multi-domain monitoring. In our evaluation, the confidentiality constraint has been relaxed, in order to investigate other performance metrics; namely, the monitoring cost, the quality of monitored paths, the anomaly detection delays, and the fairness of monitoring load distribution among domains. Simulation results on random topologies show that per-domain monitoring outperforms global monitoring for all these metrics, except the monitoring cost that is slightly lower for global monitoring.

I. INTRODUCTION

Link-level anomaly monitoring is a means to detect and localize anomalies on links that occur due to physical failures, network congestion and security attacks. It evaluates the delays and loss rates on links and their available bandwidth capacities, in order to guarantee the QoS requirements for services crossing the network. Most existing studies on link-level network monitoring have focused on mono-domain networks (e.g. [3], [4], [5], [6], [7], [8]). However, usually, services cross multiple domains that belong to different administrative authorities, and that are likely to have conflict of interests. This raises some confidentiality problems that constrain the monitoring task. Namely, most proposed monitoring schemes, which assume a detailed knowledge of the network topology, cannot be applied on multi-domain networks. This is because domains are usually not willing to disclose detailed information of their network topology and available resources.

In this paper, we focus on the problem of diagnosing link-level anomalies in multi-domain networks. This includes locating monitors and selecting monitoring paths than can cover all the links. Our goal is to come up with a monitoring

scheme that overcomes confidentiality limitations. To this end, we investigate the problem along two axes. The first axis ignores confidentiality constraints and considers the multi-domain network as a single domain. This is the global monitoring technique. The second axis overcomes confidentiality considerations by minimizing the information that is to be exchanged between domains. Each domain monitors its intra-domain links independently from the other domains, *i.e.* without disclosing any information of its intra-domain topology. Neighboring domains exchange only the set of their border nodes that are candidate to hold monitoring devices, in order to compute monitor locations and paths that can cover the inter-domain links connecting them. This is the per-domain monitoring technique. Practically, the global monitoring technique might be infeasible. However, a comparative study of these two monitoring techniques aims at finding out and evaluating all the constraints, other than confidentiality, that the multi-domain monitoring must comply to.

The problem of monitor location and anomaly detection has gained great interest over the few last years. The shared goal of all these works is to minimize the monitoring cost, that is the cost of locating monitoring devices and the cost of injecting monitoring flows along monitoring paths. In this work, we adopt the monitoring scheme proposed in [1][2] that has been proven to achieve a smooth trade-off between monitor location cost and anomaly detection cost by minimizing the two costs jointly. The main challenge is to adapt the proposed solutions to multi-domain networks with respect to topology characteristics. We provide a mathematical formulation of the problem, and we show that it is NP-Hard. Therefore, we devise a heuristic solution that takes into considerations the characteristics and the limitations of multi-domain network topology.

The proposed heuristic is used by the global monitoring technique and the per-domain technique. Besides the computation time and the monitoring cost, we consider new criteria to evaluate the two monitoring techniques. These criteria emerge from the characteristics of multi-domain networks and impact the monitoring performance. First multi-domain networks are large networks. Therefore, the global monitoring technique that considers the multi-domain network as a single domain is likely to monitor long paths that cross multiple domains. This would result in large detection delays. Indeed, the longer the monitored paths are, the larger the anomaly detection

delays are. Furthermore, long monitoring paths result in large number of suspect links in case of failure. This is because all the links of a monitoring path that exhibits an anomaly are suspect to be anomalous. Second, multi-domain networks are composed of domains that belong to different administrative and economic authorities. Therefore, the monitoring solution should distribute the monitoring load among domains fairly. Otherwise, the most overloaded domains would not be willing to collaborate.

We show through simulations that confidentiality is not the only limitation to global multi-domain monitoring. Indeed, results show that this monitoring technique delivers solutions with relatively long monitored paths and does not guarantee a fair distribution of monitoring load among domains. Besides, the computation time for the global monitoring technique is drastically high compared to the computation time for the per-domain monitoring technique. In contrast, the difference of costs of the solutions delivered by the two monitoring techniques, in terms of number of monitors and redundant measurements of links, is tiny. This is due to the characteristics of multi-domain topology that will be discussed throughout this paper.

II. PROBLEM FORMULATION

A. Network Model

We can model a multi-domain network composed of M connected domains as a set of undirected graphs $G_i = (V_i, L_i), i = 1, 2, \dots, M$. V_i is the set of nodes of domain i . It is composed of two sets: V_i^{inter} and V_i^{intra} . V_i^{inter} represents the set of border nodes that connect domain i to its neighboring domains, and V_i^{intra} represents the set of core nodes. Similarly, the set of links L_i is composed of two sets: L_i^{intra} and L_i^{inter} . L_i^{intra} represents the set of intra-domain links that connect the core nodes, and L_i^{inter} represents the set of inter-domain links that connect nodes of V_i^{inter} to the border nodes of neighboring domains. We denote by $P_i, i = 1, 2, \dots, M$ the set of intra-domain paths of domain i . A path $p \in P$ is a set of undirected intra-domain links. We denote by P^{inter} the set of inter-domain paths of the multi-domain network. A path $p^{inter} \in P^{inter}$ includes at least one inter-domain link. We refer to $G_i^{intra} = (V_i, L_i^{intra})$ as the intra-domain graph of domain i . Let $Nd = \{(i, j); i, j = 1, 2, \dots, M; i \text{ and } j \text{ are neighbor domains}\}$ be the set of neighbor domains. We refer to $G_{(i,j)} = (V_{i,j}, L_{i,j})$ as the graph of the inter-domain topology connecting domain i to domain j . $V_{i,j}$ is the set of border nodes of domains i and j that are connected to each other, and $L_{i,j}$ is the set of inter-domain links connecting domain i to domain j .

B. Problem Definition

This work addresses the problem of monitor location and network anomaly detection in multi-domain networks. For mono-domain networks, minimizing the cost of monitor location and network anomaly detection consists in deploying as few monitors as possible and avoiding redundant measurements of links, *i.e.* avoiding overlaps among monitoring paths.

These two minimization objectives are conflictual. It turned out that a joint optimization of monitor location and anomaly detection costs balances efficiently the trade-off and reduces the two costs [1]. However the problem is NP-Hard. Heuristics have been proposed for mono-domain networks in [2]. For multi-domain networks, the problem can be formulated as follows. We want to deploy monitors and select monitoring paths between the deployed monitors. The aim is to cover all the inter-domain and the intra-domain links, while reducing the number of deployed monitors and avoiding redundant measurements.

Until now, multi-domain monitoring seems to be similar to mono-domain monitoring. However, there are some constraints to multi-domain monitoring that stem from the characteristics of multi-domain networks. The first constraint is related to the structure of multi-domain networks. A multi-domain network is a set of domains that belong to different administrative authorities. Due to economic and security considerations, domains are usually not willing to share detailed information of their network topologies and resources. This is a blocking constraint to the global monitoring technique. This technique assumes the existence of a central entity that has a detailed knowledge of the intra-domain topologies of all the domains composing the multi-domain network as well as the inter-domain topologies connecting neighboring domains. An alternative solution would be to let each domain cover its intra-domain links using intra-domain paths only. Neighboring domains collaborate to cover inter-domain links connecting them. This is the per-domain technique.

At first glance, when the global topology is known, we tend to assert that the global monitoring technique outperforms the per-domain monitoring technique. This is because, considering only the metrics of number of monitors and number of redundant measurements of links, all the solutions of the per-domain monitoring technique are feasible solutions of the global monitoring technique. We illustrate our assertion in Fig. 1. and Fig. 2. Hereby, we consider a multi-domain network composed of two domains connected by a single inter-domain link (3, 5). We assume that the cost of deploying a monitor equals the cost of a redundant measurement, *i.e.* the cost of measuring a link that is already measured. Grey nodes hold monitoring devices. The black thick lines draw the monitored paths. Fig. 1. depicts a minimal per-domain monitoring solution, whereas Fig. 2. depicts a minimal global monitoring solution. We notice that the per-domain solution deploys 4 monitors, against 2 monitors and 1 redundant measurement for the global solution. The global monitoring technique succeeded to reduce the monitoring cost by removing monitors that are deployed on the border nodes of each domain.

The question that arises here is the following: *how worse is the performance of the per-domain technique compared to the global monitoring technique ?* To answer this question, we investigate the quality of the global monitoring solutions. Reducing the number of monitors results in longer monitoring paths. The figures above validate this claim. Nonetheless,




Fig. 1. Per-domain Monitoring Solution




Fig. 2. Global Monitoring Solution

multi-domain networks are usually very large networks. Subsequently, the global monitoring technique is likely to select very long monitoring paths. This is the second constraint to global monitoring; because the longer the monitored paths are, the larger the anomaly detection delays are and the larger the number of suspect links in case an anomaly occurs is. Furthermore, when domains accept to collaborate to perform global monitoring, they expect to achieve individual benefits in return. This means that the monitoring solution should distribute the monitoring load among the participating domains evenly. Therefore, besides the monitoring cost, the quality of monitoring paths and the fairness of monitoring load distribution must be considered in the evaluation of the two monitoring techniques.

Based on this discussion, we claim that confidentiality is so far not the only constraint to global monitoring, and that per-domain monitoring might turn out to be more efficient with respect to some metrics. We validate our claims in the remainder of this paper.

C. Architecture and Cost Model of Multi-Domain Monitoring

Fig. 3. depicts a sample multi-domain monitoring architecture, only nodes that hold monitoring devices are drawn. In each domain there is a Network Operation Center, denoted *Domain NOC*, that communicates with the monitors of the domain, in order to collect monitoring information and manage the monitoring task within the domain. A *domain NOC* has a detailed knowledge of the domain topology and resources. In addition, there is a central *NOC* that communicates with all the *Domain NOCs*. It collects and analyzes monitoring information collected within domains. This multi-domain architecture matches the usual architecture proposed in most works on multi-domain monitoring (e.g. [10], [11]). For global monitoring, the central *NOC* has a detailed knowledge of the topologies and the resources of all the domains; whereas for per-domain monitoring, it does not participate in the monitoring task.

We note that monitor location and anomaly detection means deploying monitors in the network and selecting monitoring paths that can cover all the network links in order to detect potential link-level anomalies. In this work, we are not interested in the localization of anomalies.




Fig. 3. Sample Multi-domain Monitoring Architecture

TABLE I
NOTATIONS USED THROUGHOUT THE PAPER

Symbol	Definition
Z_p	A binary variable that indicates whether path p is selected to be monitored
Y_{n_i}	A binary variable that indicates whether node n_i is selected as a monitor location
C_{n_i}	The cost of deploying a monitoring device on node n
C_{l_i}	The cost of monitoring the intra-domain link l_i
$C_{l_{(i,j)}}$	The cost of monitoring the inter-domain link $l_{i,j}$
$\delta_{l_i p}$	A binary parameter that indicates whether link l_i belongs to path p
$\delta_{l_{(i,j)} p}$	A binary parameter that indicates whether the inter-domain link $l_{(i,j)}$ belongs to path p
$\delta_{n_i p}$	A binary parameter that indicates whether node n_i is an end node of path p
$DR_i(SP)$	the detection ratio of path p_i considering the set of selected monitoring paths SP
CP	The set of candidate monitoring paths
SP	The set of selected monitoring paths
SM	The set of selected monitors

A summary of the symbols used in the remainder of the paper is depicted in TABLE I. The multi-domain monitoring cost can be expressed as the sum of the following costs:

* Monitor location cost: it includes the effective cost of deploying hardware and software monitoring devices and the

cost of their maintenance. In addition, it includes the cost of communications between monitors and their corresponding *Domain NOC*. For instance, the cost of communications between a monitor and the *NOC* can be expressed as a function of the physical distance that separates them. Let us denote by C_{n_i} the cost of deploying a monitor on node n_i , the multi-domain monitor location cost can be expressed as follows:

$$\sum_{i=1,2,\dots,M, n_i \in N_i} C_{n_i} Y_{n_i} \quad (1)$$

* Anomaly detection cost: it expresses the overhead of monitoring flows on the underlying network. Each link must be monitored at least once. Redundant measurements of links are considered as monitoring overhead. Let us denote by C_{l_i} the cost of measuring link l_i . C_{l_i} must be proportional to the load of link l_i , in order to avoid multiple measurements of the most overloaded links of the network. The multi-domain anomaly detection cost can be expressed as follows:

$$\begin{aligned} & \sum_{i=1,2,\dots,M, l_i \in L_i, p \in P_i \cup P^{inter}} \delta_{l_i p} C_{l_i} Z_p + \\ & \sum_{(i,j) \in Nd, l_{(i,j)} \in L_{(i,j)}, p \in P^{inter}} \delta_{l_{(i,j)} p} C_{l_{(i,j)}} Z_p \end{aligned} \quad (2)$$

III. ILP FORMULATION

The monitoring solution aims at minimizing the monitor location cost (1), and the anomaly detection cost (2). In our previous works, we have demonstrated that there is an interplay between these two minimization objectives. However, it turned out that the joint minimization of the two objectives balances efficiently this interplay [1], [2]. Therefore, our ILP formulation minimizes jointly the monitoring costs expressed above. Let α and β be relative weights of the monitor location cost and the anomaly detection cost, respectively. The objective function reads as follows:

$$\text{minimize: } \alpha*(1) + \beta*(2)$$

All the links of the multi-domain network, *i.e.* the intra-domain and the inter-domain links, must be monitored at least once. Practically, this means that each link must belong at least to one monitored path. These link coverage constraints read as follows:

$$\begin{aligned} & \sum_{i=1,2,\dots,M, p \in P_i \cup P^{inter}} \delta_{l_i p} Z_p \geq 1; \\ & \forall i = 1, 2, \dots, M, \forall l_i \in L_i \end{aligned} \quad (3)$$

$$\begin{aligned} & \sum_{p \in P^{inter}} \delta_{l_{(i,j)} p} Z_p \geq 1; \\ & \forall (i,j) \in Nd, \forall l_{(i,j)} \in L_{(i,j)}, \end{aligned} \quad (4)$$

The end nodes of all the monitored paths must hold monitoring devices. These monitor location constraints read as follows:

$$\begin{aligned} & Y_{n_i} \geq \delta_{n_i p} Z_p, \\ & \forall i = 1, 2, \dots, M, \forall n_i \in N_i, \forall p \in P_i \cup P^{inter} \end{aligned} \quad (5)$$

The equivalent problem for mono-domain networks has been shown to be NP-Hard [1]. The multi-domain monitoring problem is reduced to the mono-domain monitoring problem, introduced in [1], for $Nd = \emptyset$ and $P^{inter} = \emptyset$. We conclude that the multi-domain monitoring problem is NP-Hard, and thus, we propose a heuristic solution in the next section.

IV. HEURISTIC SOLUTION

The heuristic solution aims at minimizing the monitor location cost and the anomaly detection cost jointly, thereby balancing the trade-off between these two minimization objectives; while considering the properties and the limitations of inter-domain networks.

Multi-domain networks are, usually, composed of dense domains interconnected by few inter-domain links. Therefore, the number of paths between two nodes belonging each to a different domain, which is proportional to the number of inter-domain links separating the two domains, is small compared to the total number of paths. In [1], we have proposed a heuristic for joint optimization of monitor location and anomaly detection in mono-domain networks. This heuristic performs an in-depth exploration of the network graph, in order to find candidate monitoring paths between two given nodes. It has been shown that this technique delivers good ¹ candidate monitoring paths in a short time [1]. However, when we ran this heuristic on multi-domain networks and mono-domain networks of the same size (*i.e.* the same number of links and the same number of nodes), we noted that the computation time of the multi-domain solution is drastically higher than the computation time of the mono-domain solution. As expected, this exponential increase of the computation time is due to the computation time of candidate monitoring paths in multi-domain networks. We consider the multi-domain network depicted in Fig. 4. to illustrate our assertions.

The network is composed of two domains. The dotted links are inter-domain links and the gray nodes hold monitoring devices. The thick black line draws an inter-domain path that starts from the monitor located in Domain 1, reaches domain 2 and then returns back to domain 1. The dotted thick line draws an intra-domain path that starts from the monitor located in Domain 1, crosses the two border nodes of Domain 1 (nodes 9 and 8), but do not reach Domain 2. We note that we avoid looping paths, *i.e.* paths that cross the same nodes multiple times. These are two examples of excluded paths: long paths

¹We refer to the work in [1] for more details on the criteria and the computation technique of candidate monitoring paths in mono-domain networks




Fig. 4. Illustrative multi-domain network

that do not end at a monitoring device and whose computation time is long. It is the existence of such bad paths that makes the computation time of inter-domain candidate monitoring paths quite long, and therefore, the heuristic proposed in [1] inappropriate for global monitoring of multi-domain networks. Further, other works on intra-domain monitoring have not provided solutions for the problem of candidate monitoring path computation (e.g. [3], [4], [5], [6], [7], [8]).

A. Candidate Monitoring Path Computation in Multi-domain networks

The solution that we propose to compute candidate monitoring paths consists in assigning a positive weight to each network link, and exploring the network links with a probability that is proportional to their weights. The underlying idea is to reduce the probability to re-explore bad sequences of links, while increasing the probability to cross inter-domain links. Initially, all the links have an equal weight. This means that links have the same probability to be added to the computed path. The computation ends when the path reaches the target node, this is a good path, or when it reaches a node whose neighboring nodes already belong to the path, this is a bad path. If the computed path is good, the weights of all its links are incremented. Since all the good paths cross inter-domain links, this will increase the probability to use those links. We resume the computation of new paths from the starting node, in order to increase the space of explored paths.

B. Greedy Monitor Location and Path Selection Algorithm

Here, we give an outline of Algorithm 1. The algorithm starts by selecting two monitor locations that have the lowest costs; or in case several couple of monitor locations have the lowest cost, it selects a couple arbitrary. Next, it computes a set of candidate paths between the selected monitor locations

Algorithm 1

```

1:  $SP = \emptyset$ 
2: Select two monitor locations  $m_1, m_2$  that have the lowest
   monitor locations costs
3: Add  $m_1$  and  $m_2$  to  $SM$ 
4:  $CP \leftarrow \{\text{candidate paths between } m_1 \text{ and } m_2\}$ 
5:  $\forall p_i \in CP, DR_i(SP) \leftarrow (\# \text{ of links covered by } p_i) /$ 
   (number of links of  $p_i$  that are covered by paths in  $SP$ )
6: while ( not all links are covered ) do
7:   find  $p_s \in CP / \forall p_i \in CP, DR_s(SP) \geq DR_i(SP)$ 
8:   if ( $DR_s(SP) == 0$ ) then
9:     Go to line 25
    /* the deployed monitors cannot cover all the network
    links*/
10:  else
11:    add  $p_s$  to  $SP$ 
12:    remove  $p_s$  from  $CP$ 
13:    update  $DR_i(SP), \forall p_i \in CP$ 
14:  end if
15: end while
16: if ( Not all links are covered ) then
17:   Go to line 25
18: else
19:   if (the cost of deploying a new monitors  $\geq$  redundant
    measurements incurred by paths in  $SP$ ) then
20:     End of the algorithm
21:   else
22:     Go to line 25
23:   end if
24: end if
25: Select a new monitor that minimize the monitor location
   cost
26: Add the new monitor to  $SM$ 
27: Clear  $CP$ 
28:  $CP \leftarrow$  candidate paths between the new monitor and the
   deployed monitors
29: Remove paths that incur redundant measurements from  $SP$ 
   and add them to  $CP$ 
30: Go to line 5

```

as described above. For each candidate path, the algorithm computes a detection ratio that expresses the ratio between the number of links that are covered by the path and the number of redundant measurements, i.e. the number of links that belong to the path and that are already covered by selected monitoring paths, i.e. paths in SP . The path that have the highest detection ratio is selected. This is because it achieves the best trade-off between the number of covered links and the number of redundant measurements. The detection ratios are updated whenever a new path is selected.

Monitoring paths are selected until all the network links are covered, or all the candidate paths have their detection ratios equal to zero. In the latter case, the deployed monitors are not sufficient to cover all the network links, therefore, a new monitor is deployed. In the first case we get a full

monitoring solution, *i.e.* full coverage of the network links. However, as said earlier, we want to find the best trade-off between the monitor location cost and the anomaly detection cost. Therefore, when the algorithm gets a full solution, it verifies whether it can diminish the anomaly detection cost by deploying new monitors. It decides to deploy a new monitor if the cost of the new monitor is lower than the anomaly detection cost of the solution computed over the previous iteration.

Now, when a new monitor is deployed, the algorithm removes all the paths that incur redundant measurements from the set of selected paths and inject them into the set of candidate paths CP , and then updates the detection ratios and selects monitoring paths with respect to their detection ratios.

V. EVALUATION

In this section we first describe the evaluation methodology, and then we present and discuss numerical results.

A. Evaluation Methodology

The aim of the evaluation is to assess the performance of per-domain monitoring versus global monitoring of multi-domain networks. To this end, we run the heuristic proposed in the previous section for these two monitoring techniques on several multi-domain network topologies generated randomly using the network generator Brite (Waxman model: $\alpha = \beta = 0.4$, random node placement) [14]. Unless mentioned, we consider the following setting to generate multi-domain topologies: the network is composed of 3 domains; a domain of 10 nodes and 31 links is connected to a domain of 15 nodes and 59 links, which is in turn connected to a domain of 10 nodes and 31 links. The number of border nodes that connect each domain to a neighboring domain ranges from 2 to 3 nodes, and the number of inter-domain links between two neighboring domains ranges from 4 to 6 links. In the remainder of the paper, we refer to this setting as the default setting. Fig. 5. depicts a sample multi-domain topology. We assume that all the network nodes are candidate to hold monitoring devices and that the cost of deploying monitors is the same for all the nodes; *i.e.* $C_{n_i} = 1, \forall n_i \in N_i \forall i = 1, 2, \dots, M$. Furthermore, we assume that the link monitoring cost is the same for all the network links; *i.e.* $C_{l_i} = 1, \forall l_i \in L_i \forall i = 1, 2, \dots, M$. We assume that $\alpha = \beta = 1$. All simulation measures are the mean over 30 simulations on randomly generated topologies. Our simulation platform is developed in C++

For global monitoring, we assume that the central *NOC*, which has a global knowledge of the multi-domain topology, runs the heuristic on the global topology including the 3 domains and the inter-domain links connecting them. For per-domain monitoring, each domain runs the heuristic on its intra-domain topology. Once all the intra-domain links are covered, neighboring domains exchange their set of border nodes that hold monitoring devices, if any, in order to cover the inter-domain links connecting them using the same heuristic on




Fig. 5. Illustrative multi-domain topology

the inter-domain topology. We note that in our simulations, if two intra-domain solutions have the same monitoring cost, we choose the solution that deploys the most monitors on its border nodes so that they can be re-used to cover inter-domain links.

B. Numerical Results

We evaluate and compare the global monitoring technique and the per-domain monitoring technique along four metrics:

1) *Monitoring Cost*: we expect that the fewer are the inter-domain links, the smaller is the difference between the costs of the solutions delivered by global monitoring and the solutions delivered by per-domain monitoring. Indeed, global monitoring reduces the monitoring cost by monitoring inter-domain paths, *i.e.* paths that cross multiple domains. This is because the monitoring of inter-domain paths requires less monitoring devices and can cover links of crossed domains and also inter-domain links. However, the number of non-overlapping inter-domain monitoring paths is proportional to the number of inter-domain links. Therefore, the global monitoring optimization gets blocked by redundant measurements of inter-domain links, and ends by deploying additional monitors to avoid overlaps among inter-domain paths.

To validate our expectations, we run the heuristics for global and per-domain monitoring on topologies with the default setting, and also on topologies for which we doubled the number of inter-domain links. Fig. 6. plots the monitoring cost, *i.e.* the number of deployed monitors and the number of redundant measurements of links, for the two monitoring techniques applied on topologies with the default setting. Fig 7 plots the same metrics for the two monitoring techniques applied on topologies for which we have doubled the number of inter-domain links.

As expected, Fig. 6. shows that the difference between the monitoring costs of the solutions delivered by the two monitoring techniques is low for the default setting. We notice




Fig. 6. Monitoring Cost: default setting




Fig. 7. Monitoring Cost: doubling inter-domain links

also that the global monitoring technique deploys few monitors than the per-domain monitoring technique, whereas the number of redundant measurements is slightly larger for global monitoring. Fig. 7. shows that, compared to the results for the default setting, the global monitoring technique deploys less monitors and achieves almost the same number of redundant measurements. In contrast, the cost of the solutions delivered by the per-domain monitoring technique has almost doubled. Clearly, the per-domain monitoring techniques needs to deploy additional monitors to cover the high number of inter-domain links.

We note that practically the number of inter-domain connections is generally small in usual networks, and thus, the default setting is more realistic [9].

2) *Computation Time*: Fig. 8. draws the average CPU computation time for global and per-domain monitoring. The figure shows that per-domain monitoring is much more faster than global monitoring. As explained earlier, this is because it takes longer time to compute candidate monitoring paths that cross multiple domains than to compute intra-domain candidate monitoring paths. However, the heuristic succeeds to deliver a solution for global monitoring in about 1200 seconds, whereas other heuristics devised for mono-domain networks have stumbled against the topology properties of multi-domain networks.

3) *Quality of monitored paths*: We categorize the monitoring paths according to their lengths, in terms of number of links, into five groups: paths of length in [1-5], paths of length in [6-10], paths of length in [11-15], paths of length in [16-20], and paths of length in [21-30]. In Fig. 9., we show the distribution of network links by path length groups for the two monitoring techniques. A link belongs to a path length group if it is monitored by a path whose length is included in the length range of that group.

First, we notice that the longest monitoring paths for the per-




Fig. 8. CPU Computation Time (logarithmic scale)

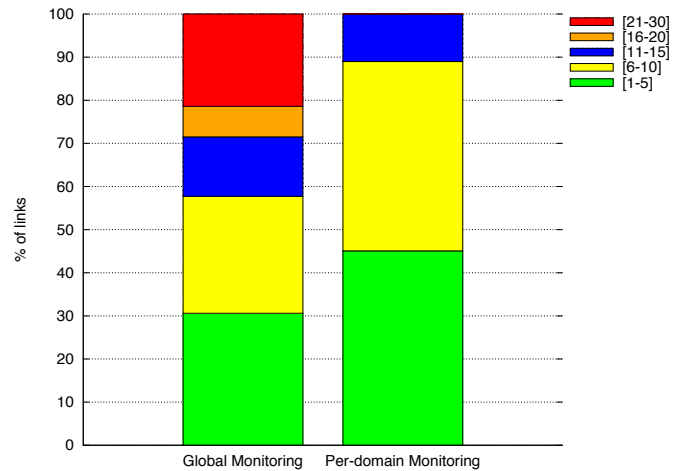


Fig. 9. Distribution of network links by path length groups

domain monitoring technique are of length less than or equal to 15 links; whereas for the global monitoring technique, the length of monitoring paths reaches 30 links. This is because the global monitoring technique monitors inter-domain paths

that are naturally longer than intra-domain paths. Second, Fig. 9. shows that more than 40% of network links belong to long paths, paths whose length is more than or equal to 15. This means that in 40% of cases of link-level anomalies, we get between 15 and 30 suspect links. In contrast, for per-domain monitoring almost 90% of network links belong to short paths, paths whose length is less than or equal to 10. We conclude that per-domain monitoring technique reduces the length of monitoring paths, and therefore, reduces anomaly detection delays and the number of suspect links when an anomaly occurs.

4) *Fairness of monitoring solutions:* In this section we propose to show the distribution of monitors and redundant measurements across domains. The aim is to evaluate the fairness of the monitoring solutions delivered by the two monitoring techniques in distributing the monitoring load among domains. To this end, we consider in our simulations multi-domain networks composed of 4 domains having the same number of intra-domain links, 18 links, and the same number of nodes, 8 nodes. Each of the 4 domains is connected to 2 other domains. The number of inter-domain connections, *i.e.* number of inter-domain links and inter-domain nodes connecting two neighboring domains, is the same for each couple of neighboring domains. For such symmetric multi-domain networks, a fair monitoring solution would distribute monitors and redundant measurements among domains evenly.

We use the Gini coefficient to measure the efficiency of monitoring cost balancing among domains [12] [13]. Fig. 10. plots the Lorenz curves for the two monitoring techniques.




Fig. 10. Distribution of monitors and redundant measurements across domains

Here, the curves are functions of the cumulative percentage of the number of domains ordered by their monitoring cost, *i.e.* the number of monitors located in the domain and the number of redundant measurements of the domain links, on the x-axis mapped onto the corresponding cumulative percentage of their monitoring costs on the y-axis. We note that if an inter-domain link is measured multiple times, we add the cost of this redundant measurement to the monitoring costs of the

two domains it connects. If the monitoring cost is distributed among domains evenly, the Lorenz curve is a diagonal line that we call the line of equality. Uneven distributions generate curves below this line. The larger is the area between the line of equality and the Lorenz curve, the greater is the inequality in the distribution of monitoring load among domains.

Fig. 10. shows that the curve corresponding to the global monitoring technique falls below the curve corresponding to the per-domain technique. This means that per-domain technique balances the monitoring load among domains more efficiently. This is explained by the fact that, in contrast to the per-domain monitoring technique, the global monitoring technique considers the multi-domain networks as a single domain. This generates uneven distributions of the monitoring load among domains.

VI. CONCLUSION

In this paper, we investigated the problem of monitor location and anomaly detection in multi-domain networks. An ILP formulation was proposed; and a heuristic that takes into account the limitations of multi-domain topologies and the requirements of multi-domain monitoring was devised. This heuristic is used to evaluate and compare two multi-domain monitoring techniques, global monitoring and per-domain monitoring, with respect to a set of performance metrics that emerge from the properties of multi-domain networks. Simulation results show that confidentiality is so far not the only constraint to global monitoring, and demonstrate that per-domain monitoring is an efficient alternative. In our future works, we will investigate the problem of anomaly localization in mono-domain and multi-domain networks.

REFERENCES

- [1] E. Salhi, S. Lahoud, and B. Cousin, *Joint Optimization of Monitor Location and Network Anomaly Detection*, IEEE LCN, 2010.
- [2] E. Salhi, S. Lahoud, and B. Cousin, *Heuristics for Joint Optimization of Monitor Location and Network Anomaly Detection*, IEEE ICC, 2011.
- [3] Y. Zhao, Z. Zhu, Y. Chen, D. Pei, and J. Wang, *Towards efficient large-scale VPN monitoring and diagnosis under operational constraints*, IEEE INFOCOM, 2009.
- [4] P. Baford, N. Duffield, A. Ron, and J. Sommers, *Network performance anomaly detection and localization*, IEEE INFOCOM, 2009.
- [5] S. Argawal, K.V.M. Naidu, and R. Rastogi, *Diagnosing link-level anomalies using passive probes*, IEEE INFOCOM, 2007.
- [6] K.V.M Naidu, D. Panigrahi, and R. Rastogi, *Detecting anomalies using end-to-end path measurements*, IEEE INFOCOM, 2008.
- [7] Y. Bajerano, R. Rastogi, *Robust Monitoring of Link Delays and Faults in IP Networks*, IEEE INFOCOM, 2003.
- [8] K. Suh, Y. Guo, J. Kurose, D. Towsley, *Locating network monitors: Complexity, heuristics, and coverage*, Computer Communication 29, pp. 1564-1577, 2009.
- [9] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson, *Measuring ISP Topologies with Rocketfuel*, IEEE/ACM Transaction on Networking, Vol. 12, NO. 1, Feb 2004.
- [10] G. Sadasivan, N. Brownlee, B. Claise, J. Quittek *Architecture for IP Flow Information Export*, RFC 5470, 2009.
- [11] C. Schmoll, E. Boschi et al., *Final Architecture Specification*, INTERMON Deliverable 15, 2001.
- [12] T. Pitoura, P. Triantafyllou, *Distribution fairness in Internet-scale networks*, ACM Transactions on Internet Technology, Vol. 9 Issue 4, 2009.
- [13] C. Dagum, *The generation and distribution of income, the Lorenz curve and the Gini ratio*, Economic Appliquee 33, pp. 327367, 1980.
- [14] BRITE, [Online]. Available: <http://www.cs.bu.edu/brite/>.