Joint Optimization of Monitor Location and Network Anomaly Detection

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

***Abstract*—Achieving cost-effective systems for network performance monitoring has been the subject of many research works over the last few years. Most of them adopt a two-step approach. The first step assigns optimal locations to monitoring devices, whereas the second step selects a minimal set of paths to be monitored. However, such an approach does not consider the trade-off between the optimization objectives of each step, and hence may lead to sub-optimal usage of network resources and biased measurements.**

In this paper, we propose to evaluate and reduce this trade-off. Toward this end, we come up with two ILP formulations for a novel monitoring cost model. The aim is to jointly minimize the monitor location cost and the anomaly detection cost, thereby obtaining a monitoring solution that minimizes the total monitoring cost. Our formulations apply for both active and passive monitoring architectures. We show that the problem is NP-hard by mapping it to the uncapacitated facility location problem. Simulation results illustrate the interplay between the optimization objectives and evaluate the quality of the obtained monitoring solution.

I. Introduction

Monitoring of network performance requires deploying monitoring devices at strategic locations to perform end-to-end path measurements. Measurements collected by the monitoring devices are usually shipped to a *Network Operations Center* (NOC) for analysis toward inferring link-level anomalies, *e.g.* loss, delay, etc. Monitoring devices can be either active beacons that simulate synthetic traffic flows or passive taps that observe real traffic flows. For active monitoring, the beacons periodically inject monitoring flows along the monitored paths. One of the most challenging tasks of monitoring systems is to come up with cost-effective monitoring schemes that efficiently detect and localize network anomalies.

Usually, the monitoring cost of such systems includes a monitor deployment cost and an operational cost. The monitor deployment cost expresses the effective cost of deploying hardware and software monitoring devices. The operational cost quantifies the overhead on the underlying network due to communications between the monitoring devices and the NOC. It also quantifies, for active monitoring, the burden on network links generated by the injected monitoring flows. Most existing works on network monitoring adopted a two-step scheme to minimize the monitoring cost: the first step, known as monitor location step, aims at minimizing the monitor deployment cost; whereas the second step, known as path selection step, aims at minimizing the operational cost.

In this paper we investigate whether the optimization objectives of the two steps are conflicting. Clearly, there is an interplay between the number of monitoring devices, the number of redundant measurements of network links, and the quality of monitored paths. On the one hand, the number of monitored paths and the number of monitoring devices should be reduced, in order to minimize the communication overhead and the monitor deployment cost. This results in monitoring few long paths that are quite likely to overlap. On the other hand, the overlaps among the monitored paths should be removed, in order to minimize the burden on network links and to avoid redundant measurements. This requires activating more monitoring devices to remove the overlaps, and results in monitoring shorter paths.

Depending on the capacities of the network, the administrator might need to assign more or less importance to the optimization of the consumption of some network ressources. For example, if the capacities of links are limited, then redundant measurements should be avoided, so that monitoring flows do not interfere with real traffic flows. However, if the optimization objectives are not properly correlated, then this would likely lead to sub-optimal usage of other resources. Therfore, it is of great importance to design a monitoring cost model that reduces the trade-off between the optimization objectives, thereby obtaining a monitoring solution that minimizes the total monitoring cost.

Our goal is to investigate and reduce this trade-off. Toward this end, we propose two different ILP formulations that model a joint optimization of monitor location and network anomaly detection problem monitoring cost model. Given a set of operational constraints and a weight for each optimization objective, our ILPs provide optimal locations for monitoring devices and optimal set of paths to be monitored that minimize the total monitoring cost and satisfies the constraints. The two ILPs were solved on randomly generated network topologies, in order to investigate the complexity of the problem and to obtain a deeper understanding of the interpaly between the optimization objectives and their impact on the quality of the monitoring solution.

II. Problem Formulation

We model the network as an undirected graph $G = (N, E)$, where N denotes the set of nodes, and E denotes the set of bidirectional edges that represent the set of links connecting the nodes. We denote by P the set of non-looping network paths, *i.e.* all paths between every pair of nodes that do not contain loops.

A solution for network performance monitoring consists of two parts: a set of locations where to deploy monitoring devices, and a set of paths that are to be monitored to detect and localize anomalies. In this paper, we are not interested on the localization of anomalies. We adopt the most common approach of anomaly detection that is monitoring a covering path set that do not distinguish link anomalies [1] (*e.g.* [1], [2], [6]). The aim of such an approach is to minimize the monitoring overhead when the network is operational. We consider a centralized monitoring infrastructure where the NOC, which has a global view of the network topology, ensures the monitor location and path selection tasks. We assume that all the network nodes are potential candidates to hold monitoring devices, and all the network paths are candidate to carry monitoring flows. However, we contend that reducing these candidate sets to sub-sets of network nodes and paths, respectively, does not impact the applicability of our model. A monitored path is defined to be a sequence of links carrying monitoring flows. The terms *monitoring device* and *monitor* are used interchangeably in the remaining of the paper. We define the monitor location cost and the anomaly detection cost as follows:

Monitor location cost: The monitor location cost expresses the effective cost of deploying and maintaining hardware and software monitoring devices. Let Cd be the cost of deploying a monitor in the network and Y_n a binary indicator if a monitoring device is located on node n , the total monitor location cost can be expressed as follows:

$$Cd \sum_{n \in N} Y_n \tag{1}$$

Anomaly detection cost: the anomaly detection cost includes two costs, a communication cost and a link measurement cost. The communication cost is the cost associated with the communications between monitors and the NOC, *e.g.* to synchronize monitors, ship measurements. Toward minimizing this cost, the monitors should be located as near as possible to the NOC. Let D_n be the distance in number of hops of node n to the NOC, the total communication cost is:

$$\sum_{n \in N} D_n Y_n \tag{2}$$

[1]The link anomalies are distinguished iff for each couple (i, j) of links there is a path p that traverses i but not j. In this case, the localization of the failed link(s) do not need additional measurements. Refer to [3] [5] for more details.

Practically, the above formulation means that the communication cost must be considered while locating monitors by privileging the nearest locations to the NOC.
The link measurement cost expresses the burden on network links due to the injected monitoring flows. Let Cl_l be the cost of injecting a monitoring flow along link l and R_l an integer counter that indicates the number of monitoring flows traversing l, the link measurement cost can be expressed as follows:

$$\sum_{l \in E} Cl_l R_l \tag{3}$$

Notice that this cost is zero for passive monitoring systems. This is because these systems do not inject synthetic monitoring flows, rather they snoop on real traffic flows. The value of Cl_l is proportional to the load of link l. The aim is to avoid redundant measurements of overloaded links.

The optimization objectives of the monitoring cost can, therefore, be summed up as follows: place as few monitors as possible at properly selected locations and reduce redundant measurements of links. We provide an example to illustrate the trade-off between these objectives. We consider two monitoring scenarios run on a network composed of 8 nodes and 10 links depicted in Fig.1 and Fig.2. In each scenario, we set the number and the positions of the monitoring devices, and then we compute the optimal set of paths to be monitored. An optimal path set must cover all the network links while minimizing redundant measurements, *i.e.* minimizes the overlaps among the monitored paths. In the first scenario (Fig.1), we locate two monitoring devices on nodes 2 and 8. An optimal path set that matches this setting is S1 = {P1(2, 5, 4, 1, 3, 6, 7, 8), P2(2, 1, 3, 6, 4, 7, 8)}. In the second scenario (Fig.1), we locate three monitoring devices on nodes 1, 6 and 8. S2 = {P3(1, 4, 6), P4(1, 3, 6, 7, 8), P5(1, 2, 5, 4, 7, 8)} is an optimal corresponding path set.

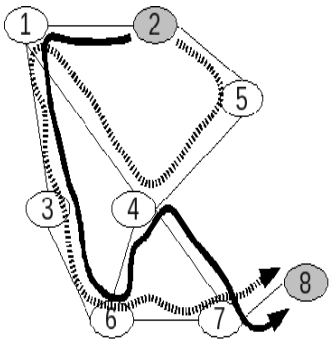

Fig. 1. Scenario 1: Two monitoring devices are deployed

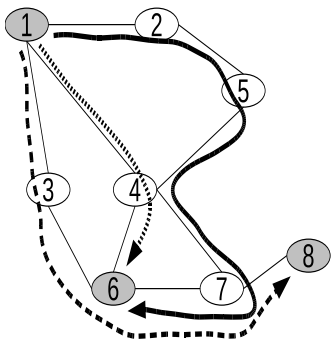

Fig. 2. Scenario 2: Three monitoring devices are deployed

In the first scenario, the links (1, 3), (3, 6) and (7, 8) are monitored twice; they belong each to two monitored paths. The deployment of an additional monitor in the second

scenario reduces the number of bi-monitored links to one link (6, 7). We conclude that the less monitors are deployed in the network, the more redundant measurements of links we obtain. In the sequel, we introduce two ILP formulations that jointly minimize the costs given by (1), (2) and (3). The first formulation is a path based formulation that takes as input the set of network paths. The second formulation is a link-flow based formulation that avoids the pre-computation of the set of network paths. The two ILPs return a set of paths that are to be monitored and a set of nodes on which to deploy monitors.

A. Path based ILP Formulation

Let us denote by Cm_n the sum of Cd and D_n. Our path based ILP formulation aims at minimizing the total monitoring cost given by the sum of (1), (2) and (3):

$$\textit{Minimize: } \alpha \sum_{l \in E, p \in P} Cl_l \delta_{lp} Z_p + \beta \sum_{n \in N} Cm_n Y_n \tag{4}$$

Z_p is a binary variable that indicates if path p is monitored, and δ_{lp} is a binary constant parameter that indicates if path p traverses link l. The number of monitoring flows traversing link l is given by the sum $\sum_{p \in P} \delta_{lp} Z_p$. α and β are positive weights that determine the relative importance of the optimization components of the above cost function.

The objective function is subject to the following constraints:

$$\sum_{p \in P} \delta_{ep} Z_p \geq 1; \quad \forall e \in E \tag{5}$$

$$Y_n \geq \delta_{np} Zp; \quad \forall n \in N, \forall p \in P \tag{6}$$

δ_{np} is a constant binary parameter that indicates if node n is an end node of path p. Constraints (5) guarantee that each network link belongs to some monitored path, whereas constraints (6) ensure that the end nodes of each monitored path are selected as monitors.

Here we show that our monitoring problem is NP-hard by mapping it to the NP-hard uncapacitated facility location (UFL) problem [10]. The UFL problem can be summed up as follows. Let F be a set of facility locations at which facilities can be built, and C a set of clients. Each facility location i is assigned a cost f_i, which represents the cost of opening facility f at location i; and each pair (i, j) of facility i and client j is assigned a client servicing cost C_{ij}, which expresses the distance between the client and the location. The client servicing costs must be symmetric and satisfy the triangle inequality, that is $C_{ij} = C_{ji}$ and $C_{ij} + C_{jk} \geq C_{ik}, \forall i, j, k \in F$. The aim of the facility location problem is to identify the locations at which the facilities shall be opened, and to serve the clients from their closest open facilities; thereby minimizing the overall cost that is the sum of the facility building cost and the client servicing cost.
Our monitoring problem can be mapped to the UFL problem as follows. The set of nodes that are candidate to hold monitoring devices maps to the set of facility locations, and the set of links maps to the set of clients. Hence, the cost of deploying a monitoring device at node n, Cm_n, matches the cost of building a facility f at a location i, f_i. Likewise, the cost of monitoring a network link l, Cl_l, matches the cost of servicing a client i from facility location j, C_{ij}. We can drop the second index related to the location of the facility in the cost of servicing a client, C_{ij}, because we assume that the cost of monitoring links is independent of the locations of the monitoring devices. This means that $C_{ij} = C_{ik}$ $\forall i, j, k \in F$, and hence the link monitoring costs are symmetric and satisfy the triangle inequality. We conclude that our monitoring problem maps to the UFL problem, and hence it is NP-hard.

B. Link-flow based ILP formulation

We expect that the path based ILP formulation would not scale to large networks where the number of paths is drastically high. In an attempt to overcome this limitation, we propose a link-flow based ILP formulation that avoids the pre-computation of the set of network paths. Beside the basic monitoring constraints, *i.e.* monitoring all the network links and selecting the end nodes of paths carrying monitoring flows as monitors, we formulate constraints that avoid forming looping paths and ensure flow conservation at nodes. We use interchangeably the terms *path* and *flow* to design a path that is candidate to carry monitoing flows.
Let $A = \{(i \rightarrow j), (j \rightarrow i); \forall (i, j) \in E\}$ be a virtual arc set, and let $Cl_{(i \rightarrow j)}$ denotes the cost of monitoring arc $(i \rightarrow j)$. We have $Cl_{(i \rightarrow j)} = Cl_{(j \rightarrow i)} = Cl_{(i,j)}$. The flows are modeled using a set of binary variables $\{X_{i \rightarrow j}(n, n'); (i \rightarrow j) \in A, (n, n') \in N^2\}$, each variable $X_{i \rightarrow j}(n, n')$ expresses whether the flow travelling between the pair of nodes (n, n') and crossing the arc $(i \rightarrow j)$ is monitored. The link-fow based ILP reads as follows:

$$\textit{Minimize: } \alpha \sum_{(i,j) \in E, (n,n') \in N^2} Cl_{(i,j)} [X_{i \rightarrow j}(n, n') + X_{j \rightarrow i}(n, n')] + \beta \sum_{n \in N} Cm_n Y_n \tag{7}$$

Subject to the following constraints:

1) Each network link must be monitored at least once:

$$\sum_{(n,n') \in N^2} X_{i \rightarrow j}(n, n') + X_{j \rightarrow i}(n, n') \geq 1; \quad \forall (i, j) \in E \tag{8}$$

2) Multiple monitoring flows might be carried between a pair of nodes [2]. We define a set of integer variables $\{W_{(n,n')}; (n, n') \in N^2\}$ to quantify the number of monitoring flows travelling between each pair of nodes. Let $IN(v)$ and $OUT(v)$ be the set of arcs entring node v and the set of arcs leaving node v, respectively. The

[2] In this case, the monitoring flows have the same end nodes, but they are carried by different paths

flow conservation constraints[3] are, hence, expressed as follows:

$$\sum_{i\to j\in OUT(v)} X_{i\to j}(n,n') - \sum_{i\to j\in IN(v)} X_{i\to j}(n,n') = \begin{cases} W_{(n,n')} & \text{iff } v=n \\ -W_{(n,n')} & \text{iff } v=n' \\ 0 & otherwise \end{cases} ; \quad \forall v,n,n' \in N \quad (9)$$

3) The following constraints ensure that the end nodes of paths carrying monitoring flows are selected as monitors:

$$Y_n \geq W_{(n,l)} + W_{(l,n)}; \quad \forall n \in N, \forall l \in E \quad (10)$$

4) Toward preventing looping flows, we define a set of integer variables $\{H_{(n,n')}(i); n,n',i \in N\}$. $H_{(n,n')}(i)$ specifies the number of hops separating node i visited by a flow travelling between the pair of nodes (n,n') from its originating node n. The idea is to force the flows to travel through nodes in an ascending order of the values of their hop variables, which prevents them from looping. We formulate the looping constraint as follows:

$$H_{(n,n')}(n) = 0; \quad \forall (n,n') \in N^2 \quad (11)$$

$$\begin{aligned} 1 - X_{i\to j}(n,n') + \frac{H_{(n,n')}(j) - 1 - H_{(n,n')}(i)}{K} \geq 0 \\ 1 - X_{j\to i}(n,n') + \frac{H_{(n,n')}(i) - 1 - H_{(n,n')}(j)}{K} \geq 0 \end{aligned} ; \quad \forall (i,j) \in E, (n,n') \in N^2 \quad (12)$$

$$H_{(n,n')}(n') \leq |N| - 1; \quad \forall (n,n') \in N^2 \quad (13)$$

Constraints (11) assign the value 0 to the hop variable of the originating node of each path, whereas constraints (13) set the upper bound of the flow lengths to the number of network nodes. Constraints (12) guarantee that flows do not re-visit an already visited node, *i.e.* a node having a value of hop variable lower than the values of those of visited nodes.

C. Extensions

The proposed ILP formulations are expressed considering active monitoring systems. The following adaptations should be introduced in order to meet passive monitoring requirements:

- The set of paths given as input in the path based ILP formulation is replaced by a set of real traffic flows that are candidate to be monitored.
- The link monitoring costs are set to 0, for passive monitoring techniques proceed by snooping on real traffic flows, and hence do not burden network links. Subsequently, the first components of our cost functions are zero.
- We might need to constrain some variables to be zero in the second formulation in order to map to real traffic flows that cross the network. For instance, if there are no flows travelling between the pair of nodes (n_i, n_j) and crossing link l_k, then $X_{l_k}(n_i,n_j)$ shall be zero.

We note that the two formulations can be easily extended to take into account other monitoring constraints. Namely, limiting the monitoring capacities of monitors (*i.e.* the maximum number of flows that can be monitored simultaneously by a given monitor), bounding the lenghts of monitored paths, bounding the number of monitored paths, etc.

III. Evaluation

In this section, we present our evaluation methodology, metrics, and simulation results.

A. Methodology and Metrics

We evaluated our ILPs using Cplex11.2 [12] running on a PC equipped with an Intel(R) Core(TM)2 Duo processor, a clock rate of 2,992.47 MHz, and 3.9 GB of RAM. All results are the mean over 20 simulations on random topologies generated using the topology generator BRITE (AS level, Waxman model) [11]. Table I depicts a summary of the main characteristics of the topologies considered in our evaluation. We devised and implemented an algorithm that computes the set of paths of an input topology. As we have anticipated owing to the complexity of the problem, we failed to compute the path set for the topologies with 12 nodes and 41 links due to memory failure. We considered an active monitoring scenario where all the network paths are candidate to be monitored and all the nodes are candidate to hold monitors, and we assumed that all the nodes are equidistant from the NOC. The values of Cl_l and Cm_n are set to 1 $\forall l \in E$ and $\forall n \in N$, respectively.

TABLE I
Summary of the topologies considered in the evaluation

Topology	# of nodes	# of links	# of paths
TOP(6, 10)	6	10	162.5
TOP(8, 18)	8	18	3176.9
TOP(10, 31)	10	31	209235.2
TOP(12, 41)	12	41	*

The aim of the evaluation is to compare the performance of the two ILP formulations, and evaluate the trade-off between the monitoring optimization objectives described in section II-B and their impact on the quality of the monitoring solution. For this purpose, we considered the following metrics:

[3]The flow that enters a node leaves it except if it is the originating node, in which case the flow only exits; or the terminating node, in which case the flow only enters

- Gap-to-optimality: it expresses the gap between the obtained solution and the optimal solution estimated by the solver. We chose to present this metric instead of the value of the objective function, because for some large topologies, the solver failed to compute an optimal solution within a reasonable time. This metric allowed us to compare the performance of the two ILPs and to validate our expectations; *(i) The path based ILP formulation is quite greedy for memory, because it must manage the set of all the network paths given as input, (ii) The link-flow based ILP formulation requires high processing capacity to handle the huge number of variables and constraints*.

- Toward studying the trade-off between minimizing the monitor cost and minimizing redundant measurements of links; we tuned the values of the weights α and β, and investigated the quality of the obtained solutions. We considered three settings: the two weights are equal, α is very large compared to β, and β is very large compared to α. Practically, the second setting matches the case where link capacities are abundant, in which case redundant measurements of links do not interfere with real traffic flows. On the opposite, the third setting matches the case where link capacities are scarce, in which case injecting additional test flows would disturb the performance of real traffic flows. For each setting, we have investigated the following metrics: the number of deployed monitors, the average length of monitored paths, the number of monitored paths and the number of redundant measurements of links.

B. Results

In the sequel, we present the simulation results. We refer to the path based ILP formulation as ILP1 and the link-flow based ILP formulation as ILP2.

1) Evaluation of the performance of the ILP formulations: In this section, we present results for $\alpha = \beta$. Tab. II presents the gap-to-optimality (GTO) and the CPU running times (RT) for the smallest topologies, *i.e.* TOP(6, 10), and the largest topologies, *i.e.* TOP(12, 41). We notice that for the smallest topologies, the two ILPs generated optimal solutions (GTO = 0%). However, the running times show that the resolution of the path based ILP is much easier than the resolution of the link-flow based ILP. This validates our assertion that the link-flow ILP is more demanding in processing capacities. This observation is confirmed in Fig. 3(a), which plots the gap-to-optimality versus the granted CPU time for the network topologies with 8 nodes and 18 links (TOP(8, 18)). Indeed, this figure shows that the path based ILP was able to obtain an optimal solution in 50.82 seconds, while after 1000 seconds, the link-flow based ILP provided a solution with nonzero gap-to-optimality. It is worth mentioning that we have granted 72 hours of CPU running time for the resolution of the link-flow based ILP on a couple of topologies, but nevertheless the solver could not provide an optimal solution.

TABLE II
EVALUATION RESULTS FOR TOP(6,10) AND TOP(12,41). GTO DENOTES GAP-TO-OPTIMALITY, RT DENOTES CPU RUNNING TIME

	ILP1		ILP2	
Topology	**GTO[%]**	**RT[sec]**	**GTO[%]**	**RT[sec]**
TOP(6, 10)	0	0.03	0	20.5
TOP(12, 41)	Out of Memory		25.01	1000

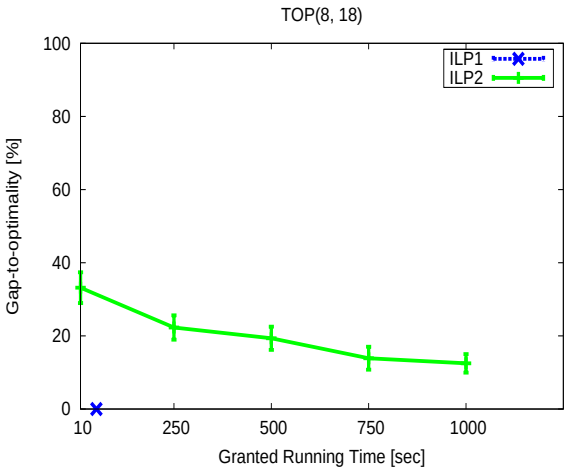

(a)

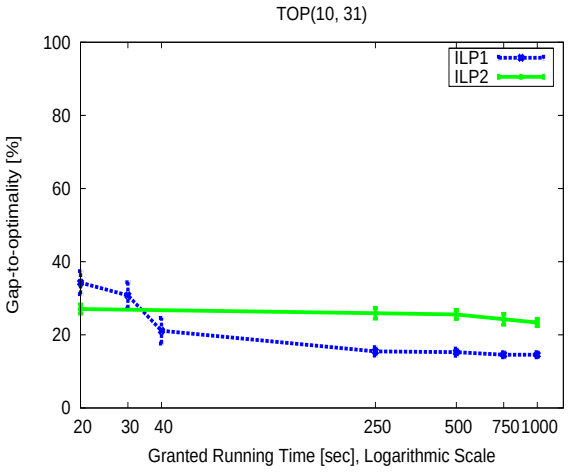

(b)

Fig. 3. Gap-to-optimality Vs. Granted Running Time

Fig. 3(b) plots the gap-to-optimality versus the granted running time for the network topologie with 10 nodes and 31 links. The results show that the two ILPs failed to generate optimal solutions within 1000 seconds. We observe that when the granted running time is small, the solutions provided by the path based ILP are worse than those provided by the link-flow based ILP; while when the granted running time

TABLE III
SIMULATION RESULTS FOR THE PERCENTAGE OF REDUNDANT MEASUREMENTS OF LINKS. % SM (SINGLE MONITORING) DENOTES THE PERCENTAGE OF LINKS MONITORED ONCE, % DM (DOUBLE MONITORING) DENOTES THE PERCENTAGE OF LINKS MONITORED TWICE

	$\beta << \alpha$		$\beta = \alpha$		$\beta >> \alpha$	
	% SM	% DM	% SM	% DM	% SM	% DM
TOP(6,10)	100	0	94	6	94	6
TOP(8,18)	100	0	90.56	9.44	90	10
TOP(10,31)	96.45	3.55	91.61	8.39	92.34	7.66

is large enough, the path based ILP performs better than the second one. This is possibly due to the large number of network paths. Indeed, the path based ILP needs to explore the input set of paths to provide a feasible solution. Compared to the results obtained by the path based ILP for the topologies with 8 nodes and 18 links, we notice that the gap-to-optimality of those obtained for the topologies with 10 nodes and 31 links goes up dramatically. This explicitly verifies that the path based ILP is quite sensitive to the size of network. The results for the topologies with 12 nodes and 41 links further demonstrates this observation. Indeed, Tab. II shows that the path flow based ILP failed to provide a feasible solution due to memory failure, while the link-fow based ILP generated a solution only 25.01% worse than the optimal within 1000 seconds.

2) Evaluation of the trade-off between the cost optimization components: now we investigate the quality of the monitoring solution versus the weight ratio $\frac{\beta}{\alpha}$ for the topologies TOP (6, 10), TOP(8,18), and TOP(10,31). As shown above, the link-flow based ILP formulation failed to generate optimal solutions within a reasonable time for the topologies with 8 nodes and 18 links, and greater. That is why, in this section, we limit our simulations on the path based ILP formulation, even though, it also failed to generate optimal solutions for the topologies with 10 nodes and 31 links. For these topologies, we show the results obtained within 1000 seconds of CPU time. We considered three settings: $\beta/\alpha = 1$; $\beta/\alpha = 10^3$, *i.e.* $\beta >> \alpha$; and $\beta/\alpha = 10^{-3}$, *i.e.* $\beta << \alpha$.

Fig.4(a) plots the average number of deployed monitors versus the network topology and the weight ratio. As expected, the figure shows that when $\beta >> \alpha$, only two beacons are deployed for all the considered topologies. This is the minimal number of monitors required to monitor a path. Obviously, the monitored paths, which have the same end nodes, are likely to overlap. This is verified in Tab.III, which shows that the percentage of redundant measurements of links ranges from 6% to 10% for the considered topologies. On another hand, Fig.4(a) shows that the number of monitors deployed when $\beta << \alpha$ is larger by several orders than those deployed for $\beta >> \alpha$, however, it is lower than to the total number of available monitor locations. This is because, the monitor cost is not zero, and hence the number of deployed monitors is also minimized in a way that minimizes the

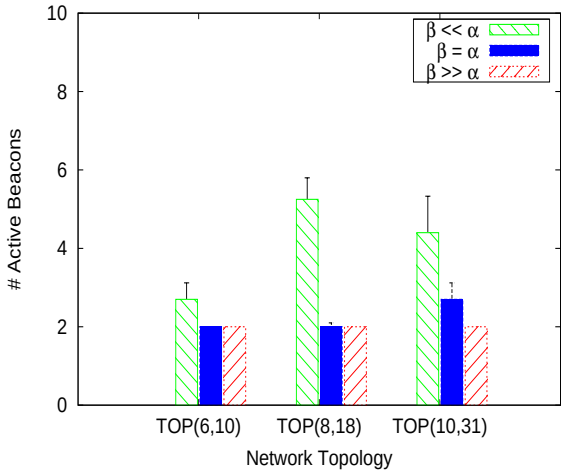

(a)

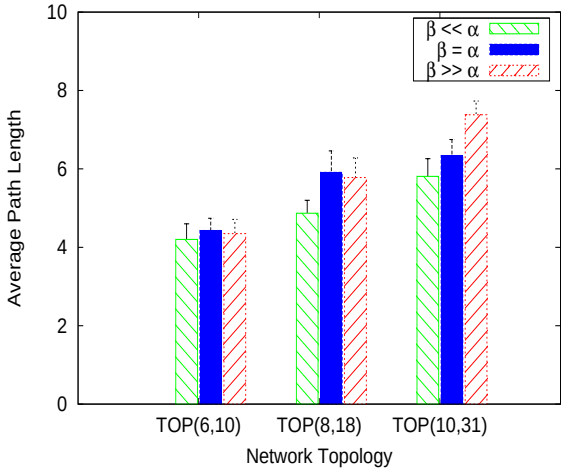

(b)

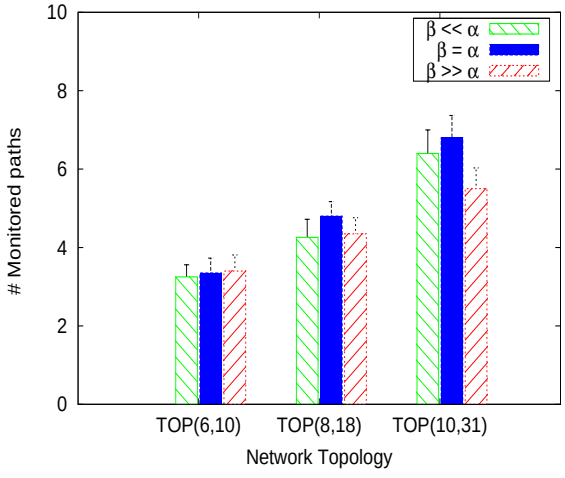

(c)

Fig. 4. Quality of the Obtained Solutions Vs. Weight Ratio and Network Topology

total monitoring cost and satisfies the monitoring constraints. Clearly, the additional monitors are deployed to remove the path overlaps. Tab.III; which shows that 100% of network links are monitored once for TOP(6,10) and TOP(8,18), whereas only 3.55% of network links are monitored twice for TOP(10,31); validates this assertion – recall that the solutions computed for the topologies TOP(10, 31) are 14.56% worse than the optimal (Fig.3(b)) –.

The above analysis results suggest that there is a trade-off between minimizing the number of deployed monitors and minimizing redundant measurements of links. However, the joint optimization of these two minimization objectives succeeds to reduce the trade-off. Indeed, Tab.III shows that less than 10% of links are monitored twice when $\beta >> \alpha$, and Fig.4(a) shows that only 60% of monitor locations are selected when $\beta << \alpha$.

Surprisingly, Fig.4(c) and Fig.4(b) show that the average number and the average lenght of monitored paths are barely sensitive to the value of the weight ratio. This meets our observation that considering the number and the length of monitored paths as the only criteria for path selection does not necessarily lead to an optimal monitoring solution. As we will show in the next section, most existing works start by locating as few monitors as possible, and then they select monitored paths. Usually the only optimization criterion of path selection is minimizing the number of paths. However, further reducing the number of paths while locating few monitors would only increase redundant measurements.

IV. Related Work

A trivial optimization of monitor selection problem consists in reducing the number of monitors toward minimizing the deployment cost, *i.e.* hardware and the software costs, and maintenance cost. Several works proposed schemes to place as few monitoring devices as possible at strategic locations of the network such that all network links are covered (*e.g.* [7]-[9]). The works in [1]-[6] addressed the minimization of the overhead of the inference techniques. Given an optimal set of monitor locations, they proposed inference schemes that monitor a small set of paths toward minimizing the communication cost. One of the most common approaches is to perform the monitoring task over two phases: anomaly detection phase and anomaly localization phase (*e.g.* [1], [2], [6]). The key goal is to reduce the monitoring overhead when the network behaves well, *i.e.* during the detection phase, by monitoring a small set of paths that covers the network links but do not distinguish anomalies. All these works decouple the monitor location problem from the path selection problem, and hence do not consider the impact of the number and the locations of monitoring devices on the quality of the monitored paths. Subsequently, none of these works reduced redundant measurements of links.

Inversely, Nguyen et al. [5] started by computing an optimal path set that enable failure monitoring, and then they computed the corresponding minimal set of monitor locations. This approach allows more flexibility in the choice of monitored paths. However, the number of required monitoring devices is not minimized when selecting the paths that are to be monitored; this may lead to deploy a suboptimal monitor set.

Recently, Zhao et al. [1] proposed an interesting formulation of the monitoring problem. Indeed, they argued that the capacities of links and monitors to handle monitoring flows should be considered while selecting monitor locations. The authors claimed that the problem is quite complex; and proposed a monitoring scheme that performs the monitoring task over multiple rounds, thereby reducing the complexity of the problem by a factor of the number of rounds. The major limitation of this multi-round approach is that it increases the delay to detect anomalies by a factor of the number of rounds.

V. Conclusion

In this paper we advocate a monitoring cost model that reduces the trade-off between minimizing the monitor location cost and minimizing the anomaly detection cost. We introduce a path based ILP formulation and a link-flow based ILP formulation, each jointly optimizes the two costs. Results show that the path based ILP is quite greedy for memory, and the link-flow based ILP is quite greedy for CPU. Hence, the two ILPs could not be used to compute monitoring solutions for large networks. However, we succeeded to validate our observations on small networks. One goal of our future work is to devise heuristics for our optimization model.

References

[1] 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.

[2] P. Baford, N. Duffield, A. Ron, and J. Sommers, "Network performance anomaly detection and localization", IEEE INFOCOM, 2009.

[3] S. Argawal, K.V.M. Naidu, and R. Rastogi, "Diagnosing link-level anomalies using passive probes", IEEE INFOCOM, 2007.

[4] K.V.M Naidu, D. Panigrahi, and R. Rastogi, "Detecting anomalies using end-to-end path measurements", IEEE INFOCOM, 2008.

[5] H.X. Nguyen, and P. Thiran, Active Measurement for multiple link failures diagnosis in IP networks, PAM Workshop, 2004.

[6] Y. Bejerano, and R. Rastogi, "Robust monitoring of link delays and faults in IP networks", IEEE INFOCOM, 2003.

[7] R. Kumar , J. Kaur, and "Efficient Beacon Placement for Network Tomography", IMC, 2004.

[8] Y. Chen, D. Bindel, and R.H. Katz, "Tomography-based overlay network monitoring", IMC, 2003.

[9] J.D Horton, and A Lopez-Ortiz, "On the number of distributed points for network tomography", IMC, 2003.

[10] G. Cornujlos, G.L. Nemhauser, and L.A Wosley, " The uncapacitated facility location problem", in P. Michandani and R. Francis, eds., Discrete Location Theory, pages 119-171, John Wiley and Sons, 1997.

[11] BRITE, [Online]. Available: http://www.cs.bu.edu/brite/

[12] Cplex, [Online]. Available: http://www.ilog.com/products/cplex.