

Active monitoring of delays with asymmetric routes

Joanna Moulierac, Miklos Molnar

► **To cite this version:**

Joanna Moulierac, Miklos Molnar. Active monitoring of delays with asymmetric routes. [Research Report] PI 1736, 2005, pp.16. inria-00000173

HAL Id: inria-00000173


<https://hal.inria.fr/inria-00000173>

Submitted on 21 Jul 2005

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

PUBLICATION
INTERNE
N° 1736



ACTIVE MONITORING OF DELAYS WITH ASYMMETRIC
ROUTES

JOANNA MOULIERAC AND MIKLOS MOLNAR

Active monitoring of delays with asymmetric routes

Joanna Moulierac^{*} and Miklos Molnar^{**}

Systèmes communicants
Projet ARMOR

Publication interne n° 1736 — Juillet 2005 — 16 pages

Abstract: There is an increasing interest in network monitoring recently. Indeed, knowledge of link characteristics is of significant importance in order to provide efficient routing. In this paper, we consider active network monitoring of link delays in a Service Provider or Enterprise IP network using round trip delays. Our proposition guarantees that all links are monitored contrary to previous propositions. Indeed, previous propositions assume symmetric routing in networks when placing the monitoring stations. With this assumption, round trips may be different when routes are asymmetric and link delays are not significant. We say that links are not monitored in this case. Previous propositions do not monitor 5.76% of links in average and 10% in worst cases during our simulations while we monitor always 100% of links. Moreover, in our proposition, the amount of traffic is reduced and the measures are more precise since the distance from a monitoring station (beacon) to the edges is limited by a given bound. Indeed, probe messages use short paths, traverse less routers and less links with our proposition. Finally, the number of beacons is not increased compared to the previous heuristic and so the installation and maintenance costs are minimized.

Key-words: active monitoring, asymmetric routing, beacons, networks

(Résumé : tsvp)

^{*} joanna.moulierac@irisa.fr

^{**} miklos.molnar@irisa.fr

La mesure active de délais de liens en considérant des routes asymétriques

Résumé : La surveillance des réseaux permet de mieux connaître les caractéristiques du réseau et de permettre par la suite un routage plus efficace. Dans cet article, nous considérons la surveillance active des délais des liens en mesurant des temps de parcours aller et retour (round trip delays). Notre proposition garantit que tous les liens sont mesurés contrairement aux propositions précédentes qui ne mesurent pas la totalité des liens du réseau. En effet, les propositions précédentes assument que le routage est symétrique et placent les stations de surveillance sous cette contrainte forte. Dans ce cas, 5.76% de liens en moyenne et 10% dans les pires cas ne sont pas surveillés pendant nos simulations alors que nous surveillons toujours 100% des liens. De plus, la quantité de trafic généré par la surveillance est réduite car nous limitons la distance d'une station au lien qu'elle surveille.

Mots clés : Surveillance active, routage asymétrique, stations de surveillance, réseaux

1 Introduction

There is an increasing interest in network monitoring recently. Network monitoring consists in collecting information about network state in order to manage resources efficiently and to ensure effective routing. Many approaches are possible in order to monitor the networks: the two most common are the passive approach and the active approach. The passive approach consists in placing devices that monitor traffic that goes through a link. This approach is often used for tomography and for networks security. The main drawback of this technique is the expensive cost of devices. Moreover, there is a technological limitation for these devices that must capture very fast information. The active approach consists, for a set of measurement points called the beacons, in sending some packets, also called probe messages, in order to detect link failures or link properties such as delays or available bandwidth. Let us notice that active monitoring involves some traffic induced by the probe messages traversing the network. One of the goal of the active monitoring design is to minimize the amount of traffic traversing the network together with the number of beacons. Indeed, the beacons induce a cost for the installation and the maintenance.

Several solutions have been proposed in the literature for active monitoring using round trip delays. The authors in [3] measure the bandwidth and latency of links with a single point of control using explicitly-routed IP packets. [9] develops a method for locating multiple link failures using active monitoring. In [6], a distributed sets of beacons on a network under BGP-like routing policy is deployed. The authors in [11] use round trip measurement and eliminates the cooperation from receivers. In [2], heuristics to monitor link latencies and faults are presented. [4] and [8] are based on the framework proposed in [2]: solutions to deal with active monitoring in dynamic networks are proposed. These two last approaches define a notion of unavoidable edges. Edges are unavoidable for a node if they belong to any shortest path tree rooted at this node. In [4] another problem is described whose goal is to choose among all the possible shortest paths trees the better routing tree for each beacon.

However, all these studies assume that routing is symmetric when placing the monitoring stations. With this assumption, some links are not monitored when routes are asymmetric as round trips may be different. Note that 30% of routes are not symmetric in Internet [10]. This problem of asymmetric routes is detailed in [1]. Last but not least, in all these studies, a beacon can monitor edges very far. This reduces the reliability of the measures because many routers are traversed by the probe messages. Moreover, this increases the traffic as many links are traversed. We assume that limiting the distance of a monitorable edge can not only increase the reliability of the measures but can also decrease the amount of traffic induced by the monitoring.

In this paper, we study active monitoring of link delays in a Service Provider or Enterprise IP network using round trip delays while taking into consideration that some routes are asymmetric in real networks.

Our main contribution are as follows. (1) We take into consideration the asymmetric routes to compute the set of monitorable edges for each beacon. By this way, our method guarantees that all the links are monitored. Previous methods do not provide this guaranty. (2) The measures of link delays are more precise and the amount of traffic is reduced since the

maximum distance from a link to its beacon is limited by a given bound. The probe messages use short paths, traverse less routers and less links. (3) Finally, the number of beacons is not increased compared to the previous algorithm in order to minimize the installation and maintenance costs of monitoring stations.

In Section 2, we present active monitoring of link delays. In Section 3, we present our proposition that guarantees that all the edges are monitored. In Section 4, we present the results of the simulations while comparing our proposition to the previously proposed algorithm in [2].

2 Active monitoring of link delays

2.1 The model

We denote the network topology as a connected graph $G = (V, E)$ where V is the set of nodes of the network and E is the set of edges. Each node $v \in V$ is assigned a set E_v of monitorable edges by v . This set E_v is deduced from the routing tree T_v of each node. The sets of the routing trees T_v can be determined by querying routing table of nodes as mentioned in [2]. The monitoring nodes, the *beacons*, can deduce properties of the links of the network G by sending *probe messages*. A centralized entity is in charge of computing the sets of beacons and to assign edges to beacons. The main problems of active monitoring are to minimize the number of beacons (in order to minimize the installation and the maintenance costs) and to minimize the cost of the probe messages (in order to minimize the amount of traffic).

2.2 How a beacon determines link delays?

The beacon in charge of monitoring a link $e = (x, y)$ sends two nearly simultaneous ICMP echo request (probe messages) to x and y at time t_x and t_y . One of these probe messages traverses e . The two extremities x and y of e answer to the beacon by sending ICMP echo reply messages. The beacon receives these messages at time t'_x and t'_y and determines the delay of the link e by making the difference between the round trip delays of the two probe messages (link e is traversed twice by one probe): $\text{delay}(x, y) = \frac{|(t'_y - t_y) - (t'_x - t_x)|}{2}$. Obviously, if there are multiple shortest paths, round trips may be different, see Fig. 1.

2.3 Active monitoring in a two-phased approach

Active monitoring algorithms are usually defined in a two-phased approach in order to minimize the number of beacons and the amount of traffic induce by the monitoring. First a set of beacons is selected with the aim of monitoring all the links of the network (beacon selection problem) and then each edge is assigned a beacon in the set of beacons which is in charge of monitoring the link (probe selection problem).

The sets E_v of monitorable edges determines which edges a node v is able to monitor. First, the goal of the Beacon selection problem is to find the minimum subset $S \subset V$ of beacons that monitor all the links:

Problem 1 *Beacon selection problem:* Given a graph $G = (V, E)$ and a set $E_v \subseteq E$ of monitorable edges for each $v \in V$, find the smallest subset of beacons $S \subset V$ s.t. $\bigcup_{v \in S} E_v = E$.

This problem is NP-Hard and a reduction from the set cover problem can be done. Second, the goal of the probe selection problem is to assign to each edge e a beacon b of S while minimizing the total cost of the probe messages. This problem is NP-Hard and an algorithm is described in section 3.3.

Problem 2 *Probe selection problem:* Given a graph $G = (V, E)$, a set $S \subset V$ of beacons and a set $E_v \subseteq E$ of monitorable edges for each $v \in V$, find for each edge $e \in E$ a beacon b_e such that $b_e \in S$ and $e \in E_{b_e}$ and while minimizing the cost of the probe messages.

The cost $c(\text{probe}(b, x))$ of a probe message $\text{probe}(b, x)$ from b to x can be equal to 1 (when minimizing the number of probes sent) or to the hop distance between b and x (when minimizing the number of links traversed by the probes).

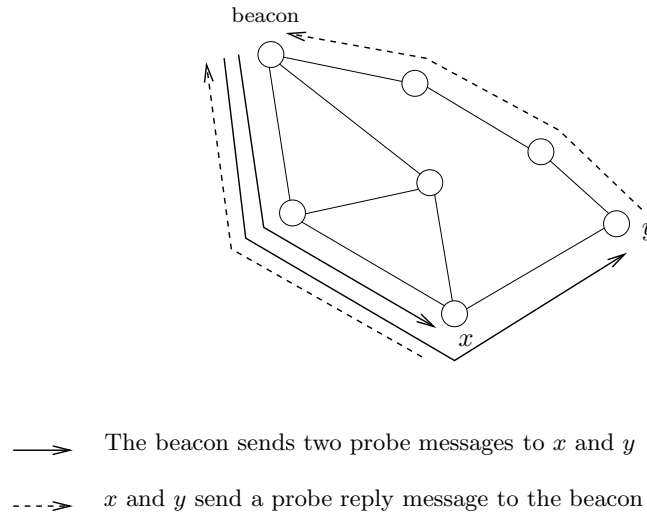


Figure 1: The probe reply messages do not follow the same path.

3 Active monitoring of link delays with asymmetric routes and limited distance

In order to choose the minimal set of beacons to monitor all the edges, we determine the set of monitorable edges for each node. Indeed, each router can monitor only a subset of links. Then, we propose a strategy to choose efficiently the beacons and finally, we select for each edge a beacon among the beacons chosen previously.

3.1 Set of monitorable edges

We define in this section how to determine the set E_v of monitorable edges for each node $v \in V$ of the network $G = (V, E)$ when considering asymmetric routes.

3.1.1 Which edges are monitorable?

Previous studies determine different ways to compute the sets of the monitorable edges for a node. For example, [2] propose to define the sets of monitorable edges E_v for a node v as the routing tree of v . [4, 8] propose to define E_v as the unavoidable edges which belong to any shortest path tree rooted at v .


We assume that the first proposition do not guarantee that all links are monitored and that the second proposition is too restrictive as the sets E_v contain few edges (only the incident edges and the edges belonging to unique shortest paths from v to other nodes of networks).

For example, on Fig. 1, [2] finds that edge (x, y) is monitorable by the beacon whereas it is not. Indeed, y sends the probe reply message using a different path than x : the two round trips are different. On bottom of Fig. 2, [4, 8] find that edge (x, y) is not monitorable by the beacon because there are multiple shortest path between the beacon and x and y and that (x, y) is not use by all the shortest paths. However, (x, y) is monitorable by the beacon in this example as round trips are equivalent.


Definition 1 *A link $e = (x, y)$ is monitorable by a node v if (x, y) belongs to the routing tree T_v of the v and if the probe reply messages sent by x and y follow the same path excluding e .*

Note that if the outgoing interface of y (the node further from v) for destination v is x , *i.e.* y uses link (x, y) to go to v , the probe reply messages follow the same path.

We do not assume by this way that routes are symmetric. Indeed, on top of Fig. 2, the routes are not symmetric ($\text{path}(T_v, v, x)$ is different from $\text{path}(T_x, x, v)$), but the link is nevertheless monitored. So, this second condition is less restricted than assuming that routing is symmetric and we guarantee that all the links are monitored.



The two probes follow exactly the same paths: routing is symmetric



Routing is asymmetric but the link is nevertheless monitored

—→ The beacon sends two probe messages to x and y

---→ x and y send a probe reply message to the beacon

Figure 2: The probe reply messages follow the same path and $e = (x, y)$ is monitored in the two figures.

3.1.2 Limiting the distance

In previous propositions, the distance from a beacon to an edge can be large: the reliability of the measures is decreased and the probe messages traverse a large number of links. In our proposition, the distance from a beacon to a monitorable edge is limited. This leads to two advantages: the measures are more precise and the cost of probes is reduced. Indeed, we know with precision the delay of a link with a beacon as extremity. For edges far from the beacon, there are additional queuing delays in the routers traversed by the probe messages. Moreover, as we limit the distance, less links are traversed and the amount of traffic is reduced.

When the distance of a monitorable edge is limited to 1, the problem can be reduced to the vertex set cover problem [7]. In this case, the probe messages traverse only the monitored link: the measures are more reliable and the cost of the probe messages is minimal. One drawback is the increase of the number of beacons.

Algorithm 1 describes our proposition to compute the set of monitorable edges.

Algorithm 1 Computation of monitorable edges

Input: Network $G = (V, E)$, routing tree T_v for each $v \in V$ and bound \mathcal{B}

Output: A set E_v of monitorable edges for each $v \in V$

For each $v \in V$

For each edge $e = (x, y) \in T_v$

if e is monitorable by v (Definition 1) **and if** $dist(v, e) \leq \mathcal{B}$ **then**

 add e to E_v

3.2 Beacon selection algorithm

The set of monitorable edges is computed for each node of the network by Algorithm 1. Then, at each step, a beacon is chosen until all the edges of the network are monitored (see Algorithm 2). The beacon chosen is the one covering the maximum number of not yet covered edges (with $|E_v|$ maximum) and which minimizes the cost of the probe messages. The total cost of the probe messages for a beacon b is the cost of the probe messages for all the edges in set E_b (i.e. $\sum_{e=(x,y) \in E_b} c(\text{probe}(b,x)) + c(\text{probe}(b,y))$). This greedy heuristic gives a $(\ln(\mathcal{V}) + 1)$, where \mathcal{V} is the size of the biggest subset of monitorable edges as can be deduced by [5].

3.3 Probe selection algorithm

The set of beacons S is computed by Algorithm 2. Each edge $e \in E$ is monitored by a beacon in S and for each edge $e \in E$, we choose the nearest beacon in S in order to minimize the cost of the probe messages. If two beacons b_1 and b_2 are at the same distance of e , we choose between b_1 and b_2 the one which monitors the less edges in order to balance the load of the beacons. The probe selection algorithm gives an approximation factor of 2 as shown in [2].

Algorithm 2 The beacon selection algorithm

Input: Network $G = (V, E)$

Output: A set of beacons $S \subset V$ monitoring E

$S \leftarrow \emptyset$

Compute E_v for each node $v \in V$ using Algorithm 1

while $E \neq \emptyset$ **do**

 Compute the set C of beacons with $|E_v|$ maximum.

 Among the beacons in C choose a beacon b that minimizes the total cost of probe messages

$S \leftarrow S \cup \{b\}$

$E \leftarrow E \setminus E_b$

For each $v \in V$, $E_v \leftarrow (E_v \setminus E_b)$

end while

4 Results of the simulations

We compare our proposition for active monitoring to heuristic Bejerano *et al.* presented in [2]. This latter proposition is very close to our work, however, there are three main differences between the two propositions:

First, heuristic Bejerano *et al.* assumes that routing is symmetric when placing beacons and determines the monitorable edges for a node v as the edges on the routing tree T_v (only the first condition presented in Definition 1).

Second, heuristic Bejerano *et al.* does not limit distance between edges and beacons. So, the results of the simulations do not vary in function of the bound \mathcal{B} as shown on the results.

Finally, at each step of the beacon selection algorithm, heuristic Bejerano *et al.* chooses the beacon which monitors the maximum number of edges not yet covered. When two beacons cover the same number of edges, they chose any of these two beacons whereas we chose the one minimizing the cost of the probe messages.

When the set of beacons is computed, the two propositions utilize the probe selection algorithm presented in subsection 3.3. The plots are the results of 500 simulations ran on random graphs with 100 nodes generated using the Waxman algorithms [12]. We compare the two heuristics with the number of beacons, the percent of edges not monitored, the maximum distance between a beacon and an edge and the number of overloaded links.

4.1 Number of beacons

As can be seen on Fig. 3, the larger the bound, the smaller the number of beacons for our heuristic. When the bound is equal to 1, the problem is reduced to the vertex set cover problem and every edge of the network has one of its extremity in the set of beacons. In this case, there is less traffic but more beacons: 60 beacons over the 100 nodes of the network are needed. When the bound is equal to 3, the number of the beacons is almost the same for

the two heuristics. In the tested topologies, 10 beacons are necessary to monitor 258 edges in average.




Figure 3: The number of Beacons

4.2 Percent of edges not monitored

Figure 4 shows the percentage of edges not monitored by heuristic Bejerano *et al.* and how often these cases happen. These edges do not respect the constraint of monitorability explained in Definition 1 and the beacons deduce wrong delays by making difference of two distinct round trips.

With heuristic Bejerano *et al.*, between 5% and 6% of edges were not monitored in 100 of 500 simulations. In average, 5.76% of edges were not monitored for the tested topologies and with some topologies, the percentage of edges not monitored reached 10%. Note that our heuristic monitors all the edges of the network because the sets of monitorable edges for each node is computed while considering the asymmetric routes.

4.3 Maximum distance of edges to beacon

Figure 5 shows the maximum distance from an edge to its assigned beacon. We plot the average of that maximum distance for the 500 simulations. Heuristic Bejerano *et al.* monitor edges which are at distance 4.34 in average. Some edges are at distance 7 from their beacons for some tested topologies whereas our heuristic always guarantees that this distance is below a given bound. With our heuristic, the maximum distance from an edge to its beacon is reduced from 4.34 (in average) to 3 while not increasing the number of beacons: less links and less routers are traversed: the delays are more precise and the amount of traffic is reduced.




Figure 4: Edges not monitored by heuristic Bejerano




Figure 5: The maximum distance of beacon to edge

4.4 The number of overloaded links

During the simulations, the cost of the probe messages was the hop-distance. This denotes the number of links traversed in the network by the probe messages. The larger the cost of the probe messages, the larger the number of links loaded by monitoring and the larger the traffic. For our heuristic, the number of overloaded links depends on the bound as shown on Fig. 6. With a bound equals to 1, the amount of traffic is minimum, but the number of beacons is maximum. With a bound of 3, our heuristic finds almost the same number of beacons as Bejerano *et al.* and the number of links traversed by the probes is reduced of around 20%. Indeed, the number of links loaded by monitoring traffic for heuristic Bejerano *et al.* is around 880 and only 730 for our heuristic: 150 less links are overloaded by monitoring traffic.




Figure 6: The cost of the probe messages

4.5 Summary of the simulation results

The bound of our algorithm allows to deal with the trade-off between the number of beacons and the cost of probe messages. The smaller the bound, the larger the number of beacons and the smaller the amount traffic. The bound 3 seems to be a good trade-off, as shown on Table 1.

	Bejerano <i>et al.</i>	Our heuristic ($\mathcal{B} = 3$)
Number of beacons	8.53	10.35
Edges not monitored	5.76%	0%
Maximum distance	4.34	3
Links loaded	878.95	730.55

Table 1: Heuristic Bejerano *et al.* and our heuristic

5 Conclusion and further work

In this paper, we proposed an improved heuristic for active monitoring of link delays in a Service Provider or Enterprise IP network while considering that some routes are asymmetric. Our heuristic guarantees that all the links of the network are monitored whereas the previous heuristic, which only considers symmetric routing, does not monitor 5.76% of links in average during our simulations. With our proposition, the measures of link delays are more precise and the amount of traffic is reduced since the distance from a link to its beacon is limited. In the worst case of our simulations, this distance is reduced from 7 for previous heuristic to 3 with almost the same number of beacons. Finally, our algorithm reduces the amount of traffic by 20% as less links are traversed by the probe messages.

This work leads to many perspectives of research. In our proposition, the routing trees are given by querying the routing table of routers. One possible extension is to investigate how to compute the set of beacons when routing tables are changing. The beacon set is no more monitoring all the links as routing trees have changed. We plan to explore this problem as part of future work. Another interesting extension is to monitor only a subset of links. A Service provider may be interested to monitor 80% of the links of the network, or to monitor some links with high traffic that may be congested and not systematically monitor all the links of the network.

Contents

1	Introduction	3
2	Active monitoring of link delays	5
2.1	The model	5
2.2	How a beacon determines link delays?	5
2.3	Active monitoring in a two-phased approach	5
3	Active monitoring of link delays with asymmetric routes and limited distance	7
3.1	Set of monitorable edges	7
3.1.1	Which edges are monitorable?	7
3.1.2	Limiting the distance	9
3.2	Beacon selection algorithm	9
3.3	Probe selection algorithm	9
4	Results of the simulations	10
4.1	Number of beacons	10
4.2	Percent of edges not monitored	11
4.3	Maximum distance of edges to beacon	11
4.4	The number of overloaded links	13
4.5	Summary of the simulation results	13
5	Conclusion and further work	14

References

- [1] G. Almes, S. Kalidinki, and M. Zekauskas. A round-trip delay metric for IPPM. RFC 2681, IETF, September 1999.
- [2] Y. Bejerano and R. Rastogi. Robust monitoring of link delays and faults in IP networks. In *IEEE Infocom*, 2003.
- [3] Y. Breitbart, C.-Y. Chan, M. Garofalakis, R. Rastogi, and A. Silberschatz. Efficiently monitoring bandwidth and latency in IP networks. In *IEEE Infocom*, pages 933–942, 2000.
- [4] Y. Breitbart, F. Dragan, and H. Gobjuka. Effective network monitoring. In *ICCCN*, pages 394–399, Oct. 2004.
- [5] V. Chavatal. A greedy heuristic for the set-covering problem. *Mathematics of Operations Research*, 4(3):223–235, 1979.
- [6] J.D. Horton and A. Lopez-Hortiz. On the number of distributed measurement points of network tomography. In *Internet Measurement Conference*, pages 204–209, 2003.
- [7] R. Karp. *Complexity of Computer Computations*, chapter Reducibility among combinatorial problems, pages 85–103. Plenum Press, 1972.
- [8] R. Kumar and J. Kaur. Efficient beacon placement for network tomography. In *Internet Measurement Conference*, Oct. 2004.
- [9] H.X. Nguyen and P. Thiran. Active measurement for multiple link failures diagnosis in IP networks. In *5th International Workshop on Passive and Active Measurement (PAM)*, pages 185–194. LNCS, April 2004.
- [10] J.-J. Pansiot and D. Grad. On routes and multicast trees in the Internet. *ACM SIGCOMM Computer Communication Review*, 28(1):41–50, Jan. 1998.
- [11] Y. Tsang, Y. Mehmet, P. Barford, and R. Nowak. Network radar: tomography from round trip time measurements. In *Internet Measurement Conference*, Oct. 2004.
- [12] B.M. Waxman. Routing of multipoint connections. *IEEE Journal Selected Areas in Communications*, 6(9), December 1988.