

Work in Progress: On the Impact of Clustering on Measurement Reduction

Damien Saucez¹, Benoit Donnet¹, Olivier Bonaventure¹ *

Université catholique de Louvain - CSE Department - Belgium

Abstract. Measuring a path performance according to one or several metrics, such as delay or bandwidth, is becoming more and more popular for applications. However, constantly probing the network is not suitable. To make measurements more scalable, the notion of clustering has emerged. In this paper, we demonstrate that clustering can limit the measurement overhead in such a context without losing too much accuracy. We first explain that measurement reduction can be observed when vantage points collaborate and use clustering to estimate path performance. We then show, with real traces, how effective is the overhead reduction and what is the impact in term of measurement accuracy.

Key words: clustering, BGP, reduction, measurement

1 Introduction

During the past decade, we have seen the emergence of a set of applications requiring more and more quality of service (QoS). For instance, IPTV needs large bandwidth and delays as low as possible. Further, while previously a content was located in a single place, it is, nowadays, frequent that this content is replicated among a set of servers located anywhere on the five continents or even among the end-users themselves. Perfect examples of this are peer-to-peer (P2P) applications and FTP mirrors.

Using measurements collected at network vantage points to infer the Internet conditions is an important feature in such a context. Indeed, the path performance metrics, such as delay, bandwidth, or packet loss, collected by applications to a potentially large set of destinations might be a good indicator on which destination to select.

However, constantly probing the network leads to scalability issues. Indeed, probes injected in the network might burden the traffic. Further, if those probes come from multiple vantage points, they might appear as being a distributed denial-of-service attack. As previously mentioned by Cheswick et al., any networking measurement system must be engineered very carefully to avoid abuse [1].

* This work is supported by the European Commission-funded 223936 ECODE project. Mr. Donnet is funded by the Fonds National de la Recherche Scientifique (FNRS – Rue d’Egmont 5, 1000 Brussels). The authors would like to gratefully acknowledge Pierre Francois and Pascal Merindol for their comments.




Fig. 1. Illustration of measurements duplication and reduction

There exist several ways for reducing the amount of required measurements. A first possibility is to modify the probing technique so that it consumes less resources. Another way is to allow collaboration between probing monitors and to cluster probe targets. *Clustering* means aggregating a subset of targets into the same hat and considering a measurements towards one of the target as being representative of the whole cluster [2–4].

In this paper, we investigate the second solution (i.e., collaboration and clustering). We deeply explain how measurement reduction can be achieved through collaboration between vantage points and destination clustering. We also discuss several metrics that can be used to evaluate the performance of a cluster based measurements campaign. We further explain that a greater measurement reduction can be achieved if collaboration between measurement sources is added to clustering.

In addition, we discuss five clustering techniques that can be used to reduce the measurements impact. These clustering techniques are based on available network information.

Based on real data collected, we evaluate these clustering techniques using the metrics we propose. We show that they are reasonably accurate while allowing a strong reduction in the amount of required probes.

The remainder of this paper is organized as follows: Sec. 2 provides a theoretical background for measurement reduction through collaboration and clustering; Sec. 3 discusses five clustering techniques and positions them regarding the state of the art; Sec. 4 evaluates these clustering techniques and shows how clustering can reduce the measurement overhead; Finally, Sec. 5 concludes this paper and discusses future directions.

2 Theoretical Background

2.1 Measurement reduction

Active Internet measurements reduction is a strategic issue when large-scale measurements are required. Up to now, several solutions have been proposed for delay measurements [5–7], Internet topology discovery [8, 9], and bandwidth estimation [10, 11]. However, these techniques, despite their strong advantages

in term of measurement reduction, are complex to deploy and can require to modify the measurement technique itself.

If keeping the measurement mechanisms intact is a requirement when trying to reduce the probing impact, two solutions are imaginable: reducing the number of measurement vantage points (i.e., the *sources*) and reducing the number of measurement targets (i.e., the *destinations*). If both lead to a measurement reduction, they can also lead to an accuracy loss. A balance must thus be found between measurement reduction and accuracy.

These two solutions might be implemented through *collaboration* and *clustering*. Collaboration can help to reduce the number of sources involved in measurements while clustering is useful to reduce the number of destinations to probe.

Any two nodes a and b could collaborate if they are topologically close. For instance, if they both belong to the same campus network. Indeed, if they share the same first hops, when measuring a destination d , it is very likely that probe packets will follow, for both a and b , the same path. Or at least, they will share large path segments. Consequently, the resulting measurement will be very close and if both nodes act in isolation from each other, they duplicate unnecessarily their efforts, as illustrated in Fig. 1(a). On the contrary, if a is aware of the measurements already performed by b to d , then, a no longer needs to measure d . In other words, collaboration avoids duplication of measurements and thus reduce the measurement overhead. Collaboration is illustrated in Fig. 1(b). How measurement sources can collaborate is still an open issue (however, efforts have been made in the context of Internet topology discovery [8]) we let for further works.

On the other hand, the key idea behind clustering is to aggregate a set of nodes under the same hat and consider that all nodes within a given hat share the same arbitrary properties. A *Cluster* \mathcal{C} is defined as a set of nodes sharing the same properties. Clustering is illustrated in Fig. 1(c).

The basic assumption behind clustering is that all nodes belonging to a given cluster share the same path performances (i.e., delay, bandwidth, etc). Consequently, measuring a single point within a cluster would be sufficient. The node that is measured to estimate the cluster path performances is called the *Reference Point*. It is thus worth to notice that clustering allows to reduce the number of measurements as only the reference point has to be measured.

In the remainder of this paper, we only consider one reference point per cluster.

A given cluster is said *popular* if any destination within this cluster is often measured. We evaluate the popularity $\pi_{\mathcal{C}}$ of a given cluster \mathcal{C} by counting the number of sources sending probes towards \mathcal{C} .

We can bring together collaboration and clustering by defining $\mathcal{S}_{\mathcal{C}}$, the *Source Cluster* of \mathcal{C} . $\mathcal{S}_{\mathcal{C}}$ represents the set of sources measuring a given cluster \mathcal{C} . Remind that this makes sense only if all nodes in $\mathcal{S}_{\mathcal{C}}$ are topologically close. Collaboration and clustering together allows a greater reduction in probing effort, as depicted

in Fig. 1(d). It is worth to notice that the set of collaborating nodes is a cluster itself.

We propose a metric for evaluating the measurement reduction: the *Measurement Reduction Factor*.

Definition 1 (Measurement Reduction Factor). *The Measurement Reduction Factor ρ for clustering technique t on \mathcal{P} is:*

$$\rho = \frac{|\mathcal{P}| - \sum_{\mathcal{C}_i \in \mathcal{C}_{\mathcal{P}}} |\mathcal{R}_{\mathcal{C}_i}|}{|\mathcal{P}|} \quad (1)$$

where \mathcal{P} is the set of $\langle s, d \rangle$ nodes pairs such that s , the source, has to measure d , the destination. $\mathcal{C}_{\mathcal{P}}$ is the set of all the clusters on \mathcal{P} assuming clustering technique t . $\mathcal{R}_{\mathcal{C}}$ is the set of all the reference points of \mathcal{C} .

Positive values for ρ means that the measurement reduction technique used effectively reduces the number of measurements. On the contrary, a negative value means that more measurements have to be performed than without reduction. For instance, if $\rho = 0.5$, the number of measurements is reduced by 50%.

Finally, note that measurement reduction is achieved on \mathcal{P} if it exists a cluster \mathcal{C} such that $|\mathcal{S}_{\mathcal{C}}| > 1$ or $|\mathcal{D}_{\mathcal{C}}| > 1$ or both.

2.2 Clustering Accuracy

In a cluster \mathcal{C} , only the reference point is used to predict the destinations performance within \mathcal{C} . It would be a matter of concern if measuring a single point (or a few points) within a cluster leads to a strong measurement accuracy loss. This section defines how to estimate the accuracy of clustering techniques.

Definition 2 (Prediction Error). *For a path performance metric m (delay, bandwidth, etc), the prediction error between node i and node j is:*

$$e_{ij} = \frac{|m_{ij} - \widehat{m}_{ij}|}{m_{ij}} \quad (2)$$

where m_{ij} is the measured value for m between i and j and \widehat{m}_{ij} its predicted value.

The prediction error gives the error proportion if the performance of a node is based on the predicted value instead of the directly measured value. The predicted error is expressed in percentage. Closer to zero the predicted error, more accurate the measurement prediction. Thus, a value of zero means that the prediction is perfect, while, for instance, a value of 0.5 means that there is a difference of 50% between the predicted and the actual value.

For any clustering technique t (see Sec. 3), the predicted metric \widehat{m}_{ij} from a node i to a node j in a cluster \mathcal{C} is the metric associated to all nodes within the cluster.

The measurement error shows how the prediction fits with the reality for a given node in a cluster. However, it can be interesting to characterize the error for the whole cluster and not only a particular node within this cluster. For such an information, statistical tools like mean, percentiles or standard deviation can be applied on the set of all cluster prediction errors.

3 Clustering Techniques

In this section, we discuss five clustering techniques. These techniques offer the strong advantage of being very easy to setup as they only require simple information already, or easily, available in the network. With these techniques, clusters are a set of IP prefixes such that a simple longest prefix matching is enough to determine to which cluster a given IP address belongs.

AS Clustering: clusters are defined based on the Autonomous System (AS) membership of Internet hosts (e.g., all the nodes of AS 2611 are put in the same cluster). This is somewhat equivalent to the notion of *super-cluster* introduced by Krishnamurthy and Wang [12] for modeling the Internet topology.

Geographic Clustering: clusters are defined based on the geographic localization of Internet hosts (e.g., all the nodes near Paris are within the same cluster).

n -agnostic Clustering: clusters are defined as fixed-length IP prefixes. All the nodes sharing the same n bits prefix are put in the same cluster. The /24 division proposed by Szymaniak et al. [3] is a particular case of such a technique (i.e., 24-agnostic clustering).

BGP Clustering: clusters are built according to BGP. Every advertised BGP prefix refers to a cluster. Each IP address belongs to the cluster with the longest matched BGP prefix. The clustering in BGP prefixes was firstly proposed by Krishnamurthy and Wang [4].

n -hybrid Clustering: clusters are built according to BGP prefixes but the minimum length of the prefixes is fixed to n . When the BGP prefix length is lower than n bits, it is divided into as many $/n$ as needed instead of the single prefix length advertised in BGP. n -hybrid Clustering is thus a mix between BGP and n -agnostic Clustering. To the best of our knowledge, we are the first to propose this technique.

Except for n -agnostic Clustering, all these clustering techniques rely on a dataset defining to which cluster an IP address belongs. For Geographical Clustering, the dataset must contain a correspondence between IP addresses and a geographical position. We believe that the MaxMind database [13] is accurate enough for most of the applications, even if it has been shown that it is less accurate (in term of geolocation) than active probing [14]. In the case of AS Clustering, the dataset consists of a mapping between an Autonomous System Number (ASN) and an IP prefix. BGP feed, iPlane [15] or Cymru [16] provide this information. In both BGP and n -hybrid Clustering, IPs are mapped to their

longest matched BGP prefix. To know the advertised BGP prefixes, a BGP feed is required. It can be obtained directly through a BGP session or a looking glass (e.g., RIPE). It is worth to notice that BGP Clustering is equivalent to 0-hybrid Clustering and that, in 32-agnostic, clusters have a cardinality of one.

We propose n -hybrid Clustering to avoid very large prefixes announced by BGP (e.g., some /8). Indeed, some of them may concern hosts spread all around the world and, thus, lead to a very low measurement accuracy.

Clustering techniques for reducing the amount of required measurements have already been extensively studied by the research community. Krishnamurthy and Wang introduce the BGP Clustering in the context of web caching [4]. In addition, they propose an adaptive clustering for addresses not classified with BGP. Unfortunately, their technique is based on reverse DNS and traceroute which is not suitable as traceroute is intrusive. Instead, we map undefined IP addresses to an agnostic / n prefix. Szymaniak et al. show that latencies are globally equivalent within the same /24 prefix on the Internet [3]. Based on that, Szymaniak et al. suggest to use a /24 cluster division (equivalent to what we call 24-agnostic Clustering in this paper) to reduce the amount of required measurement. Finally, Brown proposes to use clustering for optimizing traffic from a local network to a set of IP addresses put into a given cluster [2]. Based on measurements to a set of *scanpoints* (equivalent to our Reference Points), an entity might decide to choose a particular Internet path for traffic towards the cluster.

In addition, others proposed techniques for monitoring networks and, possibly, reducing the amount of required measurements while keeping accurate the measured metrics [17–19, 7, 9]. In particular, Donnet et al. [9] impose a limit on the number of traceroute monitor that can probe a given destination. This is done by dividing the traceroute monitors into clusters, each cluster focusing on a particular portion of the traceroute destinations.

4 Evaluation

4.1 Methodology

Our evaluation relies on two datasets. The first dataset is taken from the CAIDA’s *Archipelago* measurement infrastructure [20], the *skitter* evolution. Archipelago collects traceroute and RTT information towards all routed /24. For our study, we consider data collected in August 2008 by two Archipelago monitors: **bcn-es** (Barcelona, Spain) and **san-us** (San Diego, USA). From these dataset, we extract only complete traceroutes, i.e., traceroutes terminating at the destination. For the rest of this paper, this dataset is named *Archipelago*. For the second dataset, we collected full NetFlow traffic traces on our campus network. The traces were collected at the single 1Gbps Internet connection of the campus. A total of 45.4 TB of outgoing traffic has been monitored. However, in this paper, we do not consider 0 bytes or 0 packets flows and ignore our VPN-like traffic. Thus, after filtering, the outgoing traffic represents 7.45 TB (i.e., 22.27 Mbps on average).

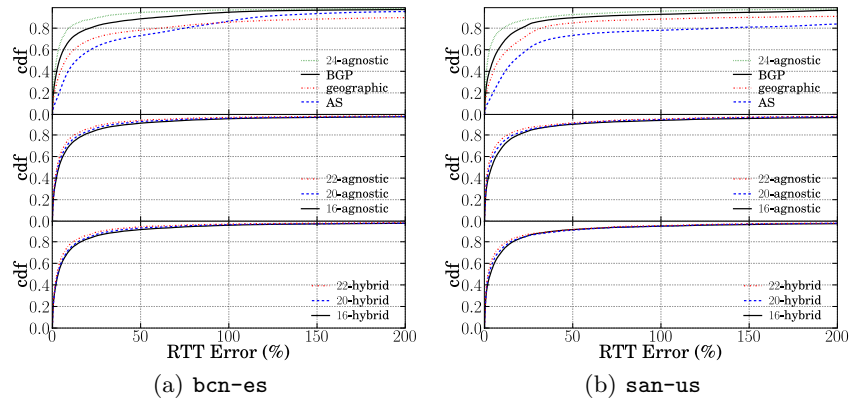


Fig. 2. Prediction error applied to RTT

The filtered dataset contains 10,084 different source IP addresses and 36,263,710 destination IP addresses for a total of 60,638,413 different layer-3 flows (i.e., the number of different $\langle \text{src}, \text{dst} \rangle$ IP address pairs). In the following of the paper this dataset is named *NetFlow*.

The NetFlow dataset is used to estimate the measurements reduction and the popularity while the Archipelago dataset permits to estimate the impact of clustering on performance estimation accuracy. Unfortunately, we have not found significant datasets that allowed us to estimate measurement accuracy and reduction at the same time. However, measurements from the **bcn-es** monitor is representative of the RTT we see on our campus network. We can thus consider the two dataset as not completely independent (**bcn-es** monitor is connected through the European research network like our campus).

The mapping between IP addresses and announced prefixes was done using the first BGP table dump of August 2008 retrieved from the University of Oregon RouteViews project [21]. ASN and IP addresses mapping was done using the iPlane dataset [15]. Finally, the geographic mapping was achieved using the MaxMind database [13].

In the following, we limit the n value to $\{16, 20, 22\}$ for n -agnostic and n -hybrid.

4.2 Measurement Accuracy

In this section, we evaluate the measurement accuracy of clustering techniques. We focus on two networking metrics: the round-trip time (RTT) and the number of hops. We let jitter and bandwidth for further work.

Fig. 2 shows the prediction error (see Def. 2) applied to the RTT for the five clustering techniques introduced in Sec. 3. Fig. 2(a) focuses on the **bcn-es** monitor and Fig. 2(b) on **san-us**. For each subfigure, there are three plots, the upper one comparing BGP, Geographic, AS, and 24-agnostic Clustering, the

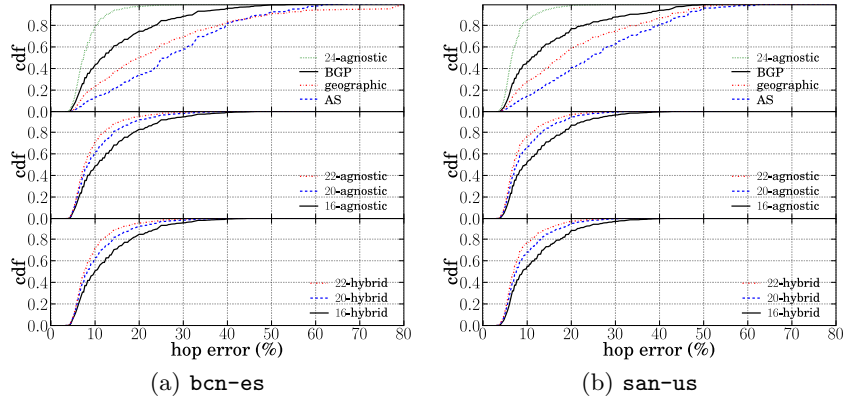


Fig. 3. Prediction error applied to hop

middle one n -agnostic, and, finally, the below one n -hybrid. The horizontal axis gives the RTT Error (in %) and the vertical axis the cumulative distribution form (CDF).

A first observation is that n -hybrid and n -agnostic Clustering perform better (n -hybrid being the best) than other clustering techniques, most of the measurements having an error less or equal to 50%. More precisely, in 80% of the cases, the RTT Error is less or equal to roughly 25%. The AS Clustering offers the worst performance. This is somewhat expected as some ASes are too large to asses the assumption that any measurement towards the cluster reference point is representative of the entire cluster. Indeed, an AS can be very large and may contain smaller entities that are separately administrated. From Fig. 2, we notice that the RTT Error might be very large, i.e., up to 1,000% for AS Clustering (not shown on Fig. 2). This is due to inherent limitations of our dataset. Some traceroute might take a long time to reach the destination and the RTT measured at this time might not be really representative of the actual “distance” separating the traceroute monitor and the destination.

Fig. 3 shows the prediction error applied to hop. Again, AS and Geographic Clustering provides the largest error. This is expected as they span larger areas. Further, for n -hybrid and n -agnostic, larger the cluster (i.e., smaller the n), larger the hop error. As for RTT error, n -hybrid provides the best performance.

Fig. 4 shows the *prediction variation* (the percentiles of all the prediction errors for a given cluster) applied to RTT (Fig. 4(a)) and to the number of hops (Fig. 4(b)). The horizontal axis gives the various clustering techniques while the horizontal axis, in log-scale, gives the prediction variation. Fig. 4 plots stacked bars. The lowest bar (i.e., the darkest) refers to the 25th percentile of the prediction variation. The middle bar shows the 50th percentile (i.e., the median), while the highest bar (i.e., the lighter) gives the 75th percentile.

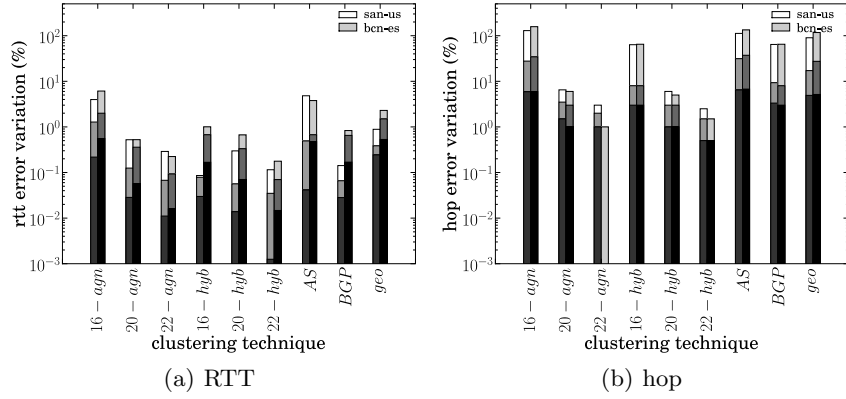


Fig. 4. 25th, 50th and 75th percentile of prediction error




Fig. 5. Popularity in the NetFlow dataset

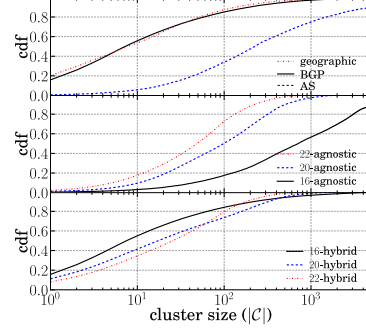


Fig. 6. Cluster size distribution

The main lesson learned from Fig. 4(a) is that, whatever the clustering technique, the RTT variation is low: the median has a maximum of 1.5% for 16-agnostic (for `san-us`). However, the situation is a little bit different for hop variation (Fig. 4(b)). The 75th percentile provides large hop errors, in particular for the largest clusters, i.e., 16-agnostic, 16-hybrid, BGP, AS, and Geographic Clustering.

To summarize, except in a few extreme cases, clustering provides thus pretty good measurement accuracy. In particular, this section suggests that n -hybrid Clustering is the most reasonable choice for a clustering technique deployment.

4.3 Measurement Reduction

In this section, we evaluate the measurement reduction observed on the NetFlow dataset for the clustering techniques introduced in Sec. 3.

We first determine if measurement reduction can be achieved, on the NetFlow dataset thanks to the collaboration. On Fig. 5, the horizontal axis, in log-scale, shows the popularity while the vertical axis gives the CDF.

Fig. 5 shows that 26.1% of the destinations are reached by, at least, 2 different sources. Collaboration will thus reduce measurements for this dataset. In addition, the top 250 destinations are reached by more than 2,000 different sources. These destinations are the major CDNs and web search engines.

Fig. 6 shows the CDF of the cluster sizes. The plot first shows that, on our dataset, clusters cover several nodes. Measurement reduction will thus be observed on our dataset.

Fig. 6 also suggests that the clustering technique influences the cluster size. For instance, 50% of the clusters cover more than 200 addresses in AS Clustering while it is the case for less than 10% in BGP or Geographic Clustering. Moreover, the cluster size is also influenced by the n parameter when applicable. With small n , prefixes are large and have the possibility to absorb many destinations, the cluster size being therefore important. As expected, n -hybrid behaves partially like BGP for small n and like n -agnostic for large n .

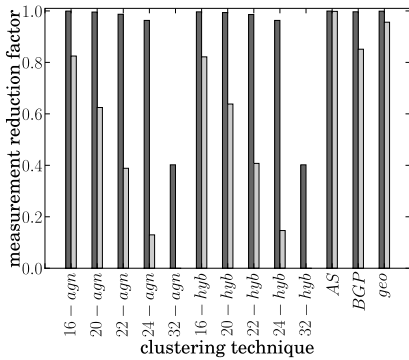


Fig. 7. Measurement reduction in the NetFlow dataset

Fig. 7 shows the effective reduction factor with respect to the different clustering techniques, with or without collaboration. The vertical axis is the reduction factor while the horizontal axis indicates the clustering techniques. On Fig. 7, light gray bars show the reduction factor when the nodes are not collaborating. On the other hand, dark gray bars show the reduction factor when nodes are collaborating.

32-agnostic and 32-hybrid Clustering (one cluster per destination) from Fig. 7 confirm that collaboration can reduce measurements (by 40% in our dataset). Light gray bars confirm that clustering can reduce measurements.

Moreover, Fig. 7 shows that, for the same clustering technique, cooperation always offers better reduction factor than no cooperation or collaboration with-

out clustering. Confirming so that combining collaboration and clustering gives even better measurements reduction than collaboration or clustering alone.

We also see on Fig. 7 that BGP, AS, and Geographic Clustering offer an important reduction factor because the number of BGP prefixes, ASes, or locations is very small compared to the number of different addresses in the NetFlow dataset.

When considering n -hybrid and n -agnostic Clustering, the reduction factor is more sensitive to n than the technique itself which is particularly true without collaboration.

Regarding to what we discussed above, we would suggest to use 20-hybrid clustering as this technique remains accurate and presents an interesting measurement reduction, even without collaboration. We would suggest not to use n -agnostic for small values of n as it will not reflect the topology. AS and geo clustering should be avoided.

In this section, we demonstrated that they effectively reduce the number of measurements while remaining quite accurate.

5 Conclusion

Measuring the quality of a set of paths, in terms of delay or bandwidth, is becoming more and more important for applications and services. Indeed, the resulting measurements could allow the application to select the best location for the required service or for getting a particular content. However, constantly monitoring the network through active measurements is not desirable due to scalability issues.

A solution for rendering measurements more scalable is to consider clustering and collaboration between measurement sources. Clustering aims at aggregating a subset of destinations into the same hat and consider that a single (or a few) measurement towards the cluster is representative of the whole cluster.

In this paper, we first discussed how collaboration and clustering might lead to a reduction in probing effort. We defined metrics for evaluating the performance of any clustering technique. Those metrics evaluates the accuracy of a prediction based on clustering and the measurement reduction.

Secondly, we explained five different clustering techniques. Those techniques have the advantage of being very easy to setup, i.e., they do not required strong calculations. Based on real data collected, we evaluated those techniques with metrics we introduced. We demonstrated that clustering techniques offer quite accurate measurement predictions as well as measurements scalability.

In this paper, we only took into account two networking metrics: delay and the number of hops. Future work should reveal how other network metrics, such as bandwidth and jitter, are impacted by cluster based measurements.

References

1. Cheswick, B., Burch, H., Branigan, S.: Mapping and visualizing the internet. In: Proc. USENIX Annual Technical Conference. (June 2000)

2. Brown, G.: Internet address space clustering for intelligent route control. see <http://www.cs.indiana.edu/~geobrown/journal-new.pdf> (2004)
3. Szymaniak, M., Presotto, D., Pierre, G., Van Steen, M.: Practical large-scale latency estimation. *Computer Networks* **52**(7) (May 2008) 1343–1364
4. Krishnamurthy, B., Wang, J.: On network-aware clustering of web clients. In: Proc. ACM SIGCOMM. (August 2000)
5. Ng, T., Zhang, H.: Predicting Internet network distance with coordinates-based approaches. In: Proc. IEEE INFOCOM. (June 2002)
6. Ng, T.S.E., Zhang, H.: A network positioning system for the Internet. In: Proc. USENIX Annual Technical Conference. (June 2004)
7. Dabek, F., Cox, R., Kaashoek, K., Morris, R.: Vivaldi, a decentralized network coordinated system. In: Proc. ACM SIGCOMM. (August 2004)
8. Donnet, B., Raoult, P., Friedman, T., Crovella, M.: Efficient algorithms for large-scale topology discovery. In: Proc. ACM SIGMETRICS. (June 2005)
9. Donnet, B., Friedman, T., Crovella, M.: Improved algorithms for network topology discovery. In: Proc. Passive and Active Measurement Workshop (PAM). (March 2005)
10. Hu, N., Steenkiste, P.: Exploiting Internet route sharing for large-scale available bandwidth estimation. In: Proc. ACM/Usenix Internet Measurement Conference (IMC). (October 2005)
11. Ramasubramanian, V., Malhki, D., Kuhn, F., Abraham, I., Balakrishnan, M., Gupta, A., Akella, A.: A unified network coordinate system for bandwidth and latency. Technical Report MSR-TR-2008-124, Microsoft Research (September 2008)
12. Krishnamurthy, B., Wang, J.: Topology modeling via cluster graphs. In: Proc. ACM SIGCOMM Workshop on Internet Measurement (IMW). (November 2001)
13. MaxMind: Geolocation and online fraud prevention from maxmind (2002) See <http://www.maxmind.com/>.
14. Siwipersad, S., Gueye, B., Uhlig, S.: Assessing the geographic resolution of exhaustive tabulation for geolocating Internet hosts. In: Proc. Passive and Active Measurement Conference (PAM). (April 2008)
15. Madhyastha, H.V., Isdal, T., Piatek, M., Dixon, C., Anderson, T., Krishnamurthy, A., Venkataramani, A.: iPlane: An information plane for distributed services. In: Proc. USENIX Symposium on Operating Systems Design and Implementation (OSDI). (November 2006) See <http://iplane.cs.washington.edu>.
16. Team Cymru: Internet security research and insight (2004) See <http://www.team-cymru.org/>.
17. Yalagandula, P., Sharma, P., Banerjee, S., Lee, S.J., Basu, S.: S3: A scalable sensing service for monitoring large networked systems. In: Proc. ACM SIGCOMM Workshop on Internet Network Measurement (INM). (September 2006)
18. Leonard, D., Loguinov, D.: Turbo King: Framework for large-scale Internet delay measurements. In: Proc. IEEE INFOCOM. (April 2008)
19. Mahajan, R., Zhang, M., Poole, L., Pai, V.: Uncovering performance differences among backbone ISPs with NetDiff. In: Proc. Symposium on Network Systems Design and Implementation (NSDI). (April 2008)
20. Hyun, Y., Huffaker, B., Andersen, D., Aben, E., Shannon, C., Luckie, M., claffy, k.: The CAIDA IPv4 Routed /24 Topology Dataset - August 2008 http://www.caida.org/data/active/ipv4_routed_24_topology_dataset.xml.
21. University of Oregon: Route views, University of Oregon Route Views project See <http://www.routeviews.org/>.