A scheme for measuring subpath available bandwidth

Andreas Johnsson, Svante Ekelin and Christofer Flinta
Ericsson Research
Kista, Sweden
{andreas.a.johnsson, svante.ekelin, christofer.flinta}@ericsson.com

Abstract—This paper presents a novel probing scheme which can be used for estimating the available bandwidth of subpaths, without the requirement of control over both endpoints of a network path. Instead of a probe-packet receiver, this scheme uses the ICMP capability of routers. An estimate of the available bandwidth from the endpoint to a router is obtained in much the same way as for state-of-the-art end-to-end probing methods. Taking into account ICMP packet generation limitations and delay, the estimate should be interpreted as a lower bound of the actual available bandwidth.

By combining estimates for several routers along a path, it is possible to obtain lower bound estimates also for subpaths between routers. These bounds may be further sharpened if combined with an estimate from an end-to-end measurement. From the obtained subpath bounds, it may be possible to identify a portion of the end-to-end path containing the bottleneck link. It is also possible to infer lower bounds for the individual router hops.

In order to demonstrate the feasibility of this ICMP-based scheme, the paper reports on a measurement study conducted over the Internet.

Keywords: *Available bandwidth, active probing, ICMP, subpath*

I. INTRODUCTION

Active measurement methods utilizing self-induced congestion [1] [2] for estimating the end-to-end available bandwidth [3] [4] [5] have had considerable attention in recent years. One example of commercial applications using active probing for measuring bandwidth is described in [6]. Probe packets are sent between Xbox 360 consoles in order to find player with good enough network performance.

The concept of self-induced congestion is to inject probe packets at a specified rate at one endpoint and then observe the changes in the rate at another endpoint. This is illustrated in Fig. 1. Probe packets are sent with an initial probe rate u. If the probe-packet rate is higher than the available bandwidth the packets will congest the path momentarily and thereby they will on average be received at a reduced rate. This is equivalent to that the time dispersion at the sender side, Δ_{out}, is larger than the time dispersion at the receiver, Δ_{in}. That is, $\varepsilon > 0$. On the other hand, if the probe-packet rate is lower than the available bandwidth, there is no congestion and so on average no time dispersion change, that is $\varepsilon = 0$.

BART [7] and pathChirp [8] are two examples of methods utilizing self-induced congestion that also are capable of providing available bandwidth estimates in real time. By sending probe packets between two endpoints the end-to-end available bandwidth is estimated by means of various statistical methods where Δ_{out} is compared to Δ_{in}.

This paper presents a novel scheme for actively measuring the available bandwidth from an endpoint to ICMP-enabled routers along a network path utilizing self-induced congestion. That is, a separate probe-packet receiver endpoint is not required.

The endpoint injects probe packets destined to another IP node, selected in such a way that the path includes the ICMP-enabled routers of interest. These probe packets are hop-limited, that is their TTL value is set to expire at a specific hop distance along the path. When the TTL value expires a probe packet is discarded. An ICMP-enabled router will respond with an ICMP time exceeded packet, which travels back to the endpoint that sent the probe packets.

The endpoint records the time both when sending the probe packets and when receiving the corresponding ICMP packets. These sending and receiving times can then be used for estimating available bandwidth of the path between the endpoint and the router discarding the probe packets. In contrast to end-to-end measurement methods, this proposed scheme combines the sender and receiver functionality in one node.

Note that we do not suggest replacing existing end-to-end measurement methods with the scheme proposed in this paper. A major advantage of this scheme is being able to measure available bandwidth without access to a remote endpoint. Further, by combining this scheme with an end-to-end estimation method, one can achieve considerably better performance than by simply employing either of them on its own.

It may then be possible to gain information on the available bandwidth for subpaths of the end-to-end path. With this information it is therefore possible to identify which part of the path that contains the bottleneck. Note that this higher resolution of the path characteristics is obtained without any need of traffic tapping nodes within the network.

The rest of the paper is organized as follows; Section II provides an overview of related work. Section III discusses definitions and notations. Section IV presents the ICMP probing scheme in detail while Section V discusses aspects of the measurement methodology. In Section VI the methodology for resolving subpath estimates is presented while Section VII provides an empirical feasibility study. Summary and conclusions are located in Section VIII.

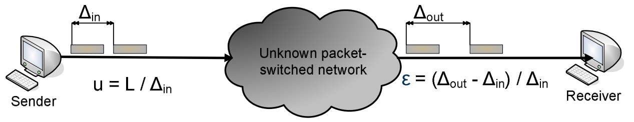

Figure 1. Illustration of end-to-end active probing to obtain available bandwidth estimates. L is the probe packet size and u is the probe-packet rate.

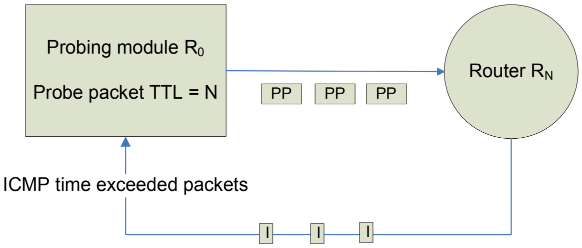

Figure 2. Illustration of how to utilize router ICMP time exceeded capabilities for measuring available bandwidth. Boxes labelled PP corresponds to UDP probe packets. These packets have a TTL = N when sent into the network. Router R_N discards the probe packets and returns ICMP time exceeded packets labelled I to R_0.

II. RELATED WORK

According to [9], bandwidth estimation techniques can be classified in various ways of which two are reviewed below. One important classification is based on whether the tools conduct hop-by-hop or end-to-end measurements of a network path.

Usually hop-by-hop techniques rely on packets with a TTL small enough in order not to reach the destination; that is, hop-limited packets. As proposed by the new scheme in this paper hop-limited packets are used to trigger ICMP responses from intermediate routers and thus information can be retrieved from each hop along a path. Hop-by-hop methods are often used for estimating link properties such as capacity for each hop along a path. Such techniques usually do not rely on self-induced congestion. Examples of methods in this category are pchar [10] and pathRate [11].

End-to-end techniques on the other hand usually utilize the concept of self-induced congestion. These methods most often estimate the end-to-end available bandwidth rather than the hop-by-hop link properties. Examples of methods in this category are BART and pathChirp.

A second classification according to the same paper [9] is based on whether the tools measure the bottleneck link capacity or the available bandwidth of a path. The bottleneck link capacity is the maximum transmission rate that could be achieved between two endpoints of a path in absence of any competing traffic. The available bandwidth on the other hand is the unused capacity that could be acquired by other network traffic flows without congesting the path.

The proposed scheme in this paper relies on self-induced congestion and hop-by-hop ICMP packet generation. In contrast to other hop-by-hop techniques, the method in this paper aims at estimating the available bandwidth, not only bottleneck or hop-by-hop link capacity. The method proposed in [9] uses a combination of hop-by-hop and end-to-end probing techniques; however, that method aims at determining the bottleneck link capacity and not the available bandwidth for each hop.

III. DEFINITIONS AND NOTATIONS

For the purpose of this paper, a path $P(R_0, R_N)$ is defined as a sequence of network nodes $R_0, \ldots R_N$. In the case of an end-to-end path R_0 correspond to the source node and R_N to the destination node. R_k corresponds to node number k on the path $P(R_0, R_N)$.

A subpath $P(R_i, R_k)$, where $0 \leq i < k \leq N$, corresponds to the part of the path $P(R_0, R_N)$ from R_i to R_k. A subpath is also a path. Conversely, $P(R_0, R_N)$ is called a superpath of $P(R_i, R_k)$.

A path $P(R_0, R_k)$ consisting of at least three network nodes can be partitioned into two subpaths $P(R_0, R_i)$ and $P(R_i, R_k)$ where $0 < i < k$. Recall that the available bandwidth of the path $P(R_0, R_k)$ is the minimum of the available bandwidths of the two subpaths. The path $P(R_0, R_k)$ is a superpath of both subpaths.

The estimate of the available bandwidth for a path $P(R_0, R_N)$ is denoted by $A(P(R_0, R_N))$.

IV. A SCHEME FOR ESTIMATING AVAILABLE BANDWIDTH EXPLOITING ICMP CAPABILITIES IN ROUTERS

This section describes a novel scheme for estimating available bandwidth of a subpath from an endpoint to an ICMP-enabled router. This is illustrated in Fig. 2 and described below.

1. The endpoint R_0 sets TTL = N for the probe packets to be sent. The destination address is selected in such a way that the path includes the ICMP-enabled routers of interest
2. The endpoint injects the probe packets with pre-determined inter-packet separation and packet sizes
3. The send time is recorded for each packet
4. For each router hop the TTL is decreased by one. The probe packet is discarded if TTL = 0, see Fig. 2. The router R_N generates an ICMP time exceeded packet for each corresponding probe packet that is discarded
5. The ICMP packet is returned to the endpoint that injected the probe packet
6. The endpoint records the arrival time of each ICMP time exceeded packet and calculates the inter-packet separation
7. The endpoint now possesses both send and receive time stamps and information about packet sizes and can thus utilize e.g. BART or pathChirp for producing estimates of $A(P(R_0, R_N))$

The algorithm can be repeated for several TTL values in order to estimate other subpaths. This enables further conclusions to be drawn regarding the available bandwidth, including resolving into subpaths $P(R_i, R_k)$.

V. ASPECTS OF MEASUREMENT METHODOLOGY

There are several aspects of the usage of the ICMP-based probing scheme presented in Section IV. These are discussed in the following subsections.

A. Matching ICMP packets to probe packets

Packet reordering and loss are problems that active measurement methods need to address. For end-to-end measurement methods it is often enough to introduce a packet identifier in the UDP data field. However, using the scheme presented in this paper the identifiers are lost when the packets are discarded. Two techniques for overcoming this problem are discussed in the following paragraphs.

ICMP time exceeded packets contain the IP header and the following eight octets of the discarded packets [12]. This includes the UDP/TCP source and destination port as well as the checksum. This can be used in several ways.

Under the assumption that the UDP checksum function is in use and that each UDP packet has a unique identifier in the data field, the checksum returned in the ICMP packet will also be unique. Thus, the checksum can be used to match UDP packets to the corresponding ICMP time exceeded packets.

Another proposal is to utilize the source or destination port information that also is embedded in the ICMP time exceeded packets. For a probe-packet train, each packet can be sent to or from a unique port [12]. Thus, matching can be made between the UDP packets and the ICMP time exceeded packets.

B. Distinguishing forward and reverse path bottlenecks

The measurement traffic travels in the form of probe packets (UDP) from the endpoint to the router where the TTL expires and as ICMP time exceeded packets on the path back to the endpoint. Therefore, it is not immediately obvious whether the available bandwidth estimate is related to a bottleneck in the direction from the endpoint to the router or vice versa.

However, using large probe packets it is in most practical cases possible to distinguish between the forward and reverse path. If congestion is detected, it is very unlikely that this is related to a bottleneck on the reverse path from R_N back to R_0. In Fig. 2 this is illustrated. The probe packets can be designed to be much larger than the ICMP time exceeded packets. For example, in the implementations of both BART and pathChirp the probe packet sizes are of the order of 1500 bytes, while the ICMP packets have an approximate size of 70 bytes. This means that for each probing sequence, the probing rate is roughly 20 times higher in the forward as compared to the reverse direction. Self-induced congestion could occur on the reverse path only if this path has an available bandwidth less than 1/20 of that of the forward path. Thus the probability of self-induced congestion is substantially lower on the reverse path.

VI. RESOLVING SUBPATH ESTIMATES

By combining measurements with different TTL values it is possible to further characterize the path. This is discussed in the next subsections.

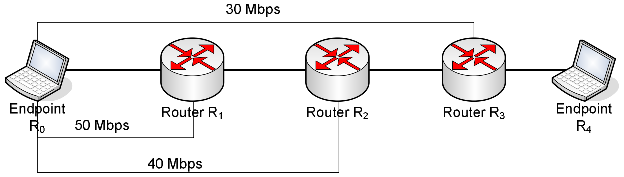

Figure 3. An example network showing available bandwidth estimates for $P(R_0, R_1)$, $P(R_0, R_2)$ and $P(R_0, R_3)$. The routers are determined by the addresses of R_0 and R_4.

A. Resolving subpath estimates

In this subsection it is assumed that the ICMP packets are generated without delay. Deviations from this case are discussed in the next subsection.

The scheme presented in this paper can be used to measure the available bandwidth of subpaths corresponding to different TTL values. For example, $P(R_0, R_i)$ corresponds to TTL = i and $P(R_0, R_k)$ to TTL = k. If $i < k$, $P(R_0, R_i)$ is a subpath of $P(R_0, R_k)$.

Now conclusions about the subpath $P(R_i, R_k)$ can be drawn, using the fact that the available bandwidth of the path $P(R_0, R_k)$ is the minimum of the available bandwidths of the two subpaths $P(R_0, R_i)$ and $P(R_i, R_k)$. This is simply due to the fact that the available bandwidth of a subpath cannot be lower than the available bandwidth of a superpath.

An example is illustrated in Fig. 3. For TTL = 1 the estimated available bandwidth of the subpath $P(R_0, R_1)$ is 50 Mbps. Increasing the TTL to 2 gives an estimate $A(P(R_0, R_2))$ = 40 Mbps. It can then be inferred that the available bandwidth $A(P(R_1, R_2))$ from router R_1 to router R_2 is 40 Mbps. For the path $P(R_0, R_2)$ this means that the bottleneck is between router R_1 and router R_2.

Exploring further, one finds that the available bandwidth estimate using TTL = 3 is 30 Mbps according to the figure. Using the same reasoning $A(P(R_2, R_3)) = 30$.

This illustrates that it may be possible to resolve which subpath contains the bottleneck as well as the available bandwidth for that subpath.

B. Implications of ICMP packet generation limitations on available bandwidth estimates

The generation of ICMP packets in a router may be affected by the forwarding and routing of packets as well as system operator policies defining how and when to generate ICMP packets. If the router load is high, it is likely that the generation of ICMP time exceeded packets is delayed. This will also contribute to the time dispersion of the ICMP packets received at the endpoint. That is, the actual available bandwidth can be higher than or equal to the produced estimate.

This means that the available bandwidth estimates obtained from utilizing the time dispersion of the ICMP packets shall be interpreted as lower bounds of the actual available bandwidth of a path. The available bandwidth estimate obtained in step 7 in the scheme in Section IV shall therefore be interpreted as a lower bound.

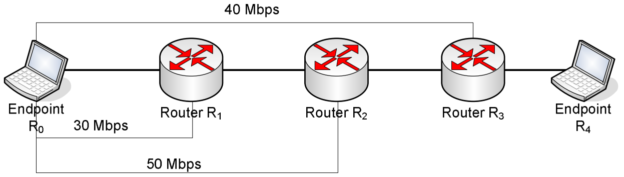

Figure 4. An example network showing available bandwidth estimates for $P(R_0, R_1)$, $P(R_0, R_2)$ and $P(R_0, R_3)$.

The above reasoning must also be applied when deducing subpath characteristics as described in Section VI.A. This means that $A(P(R_i, R_k))$ shall be interpreted as a lower bound. Consequently, conclusions about where the bottleneck is located can not be drawn directly.

An example illustrates the above discussion, Fig. 4 depicts an endpoint R_0 having probe-packet sender capabilities. The network path and hence the routers to be used is determined by the IP address of endpoint R_4. The available bandwidth estimates for several subpaths $P(R_0, R_k)$ corresponding to TTL = k are also shown.

For TTL = 1 the estimate obtained is $A(P(R_0, R_1)) = 30$ Mbps. Increasing the TTL to 2 gives an estimate of 50 Mbps. This indicates that $A(P(R_0, R_1))$ was ”artificially” reduced, that is not by real congestion but by ICMP generation limitations at R_1. This means that the available bandwidth from the endpoint to the first router is at least 50 Mbps.

Exploring further, one finds that the available bandwidth estimate $A(P(R_0, R_3)) = 40$ Mbps. This implies that $A(P(R_2, R_3)) = 40$ Mbps. Note that these estimates are interpreted as lower bounds.

The algorithm described above does not in detail cover the issue on how to characterize the path, but the reasoning gives the general idea to be used for determining a “lower bound” of the available bandwidth of subpaths.

C. *Combining the ICMP-probing scheme and end-to-end measurement methods*

By combining this ICMP-probing scheme with end-to-end measurement methods it is possible to draw further conclusions. End-to-end measurements of available bandwidth normally give the available bandwidth of the path rather than a lower bound. This means that it may be possible to draw conclusions regarding the location of the bottleneck of the path. This is further elaborated in Section VII.B.

VII. EMPIRICAL FEASIBILITY STUDY

In this section an empirical feasibility study is presented, illustrating also the potential of using this scheme in combination with traditional end-to-end methods. First, results showing the ICMP response rates from a set of routers in Sweden are presented. Then, one specific path between Ericsson Research in Stockholm and Uppsala University has been characterized by deploying the ICMP-based probing scheme together with end-to-end probing.

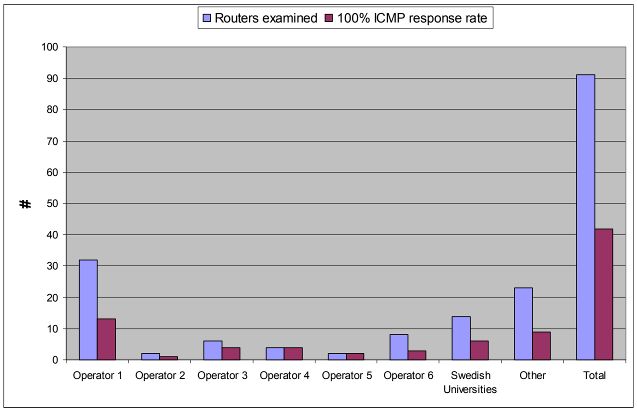

Figure 5. The total number of routers examined in comparison to the number of routers responding with the required number of ICMP time exceeded packets.

A. *Router ICMP response rate*

In order to utilize the ICMP probing scheme, at least one router must generate the required number of ICMP time exceeded packets. That is, one ICMP packet for each UDP probe packet injected into the network path. As already discussed in the paper, the more routers that could be utilized along a path the more information is gained.

Even though the present scheme requires that each discarded UDP packet generates an ICMP time exceeded packet, future improvements may be able to relax this requirement somewhat.

In Fig. 5, the number of routers responding with one ICMP time exceeded packet for each discarded UDP probe packet is compared to the total number of routers studied. The routers are operated by various operators in Sweden, as shown in the figure. For all operators studied we have encountered at least one router that fulfils the requirement. Operators 1 through 6 are commercial operators. The Swedish University routers belong to either a specific university or to SUNET (the Swedish University Network). Routers categorized as “Other” belong to unidentified operators. For each router where the TTL expires, 100 probe-packet trains, each containing 17 probe packets, are injected into the network path. One probe-packet train is injected per second. That is, the routers have to generate 1700 ICMP time exceeded packets in total during a time period of 100 seconds.

The routers have been selected by collecting traceroute information from paths between one computer at Ericsson Research and a set of university, municipality and commercial Swedish hosts covering a large part of the country. From the figure it can be concluded that about half of the routers did respond with an ICMP time exceeded packet to all discarded probe packets.

B. *Characterization of an Internet path*

For the purpose of illustration of the potential of the technique outlined in this paper, this section presents a study where a path between Ericsson Research in Stockholm and Uppsala University has been characterized in terms of lower bounds for subpath available bandwidths and of location of the bottleneck. The ICMP scheme is combined with end-to-end probing using BART [7].

The measured path consists of 10 router hops. The first two routers belong to Ericsson in Stockholm; the ownership of router R_3 is unknown; routers R_4 through R_8 are part of the network owned by a commercial operator (“Operator 1”); while router R_9 and R_{10} belong to Swedish Universities (compare Fig. 5).

For each measurement (i.e. for each TTL) 100 probe-packet trains each containing 17 probe packets of size 1500 bytes are injected once a second. Previous findings indicate that 17 probe packets in each train provide reasonably accurate available bandwidth estimates [7]. BART requires an initialization phase, that is, a sequence of measurement samples are needed for the estimates to “close in” on the true value. 100 probe-packet trains are usually enough.

Fig. 6 compares the number of probe packets injected to the number of ICMP time exceeded packets received from each router. As can be seen, five out of the ten routers reply with ICMP time exceeded packets to all ”requests” and can thus be used by the proposed scheme.

The measurements were performed in sequence for TTL = 2, 3, 6, 8, 9; each value corresponding to a specific usable router on the path.

Fig. 7 shows the available bandwidth estimates between the endpoint R_0 and routers R_2, R_3, R_6, R_8 and R_9 together with the end-to-end available bandwidth (i.e. between R_0 and R_{11}). The first samples in each measurement sequence correspond to an initialization phase.

The initialization phase for the measurement corresponding to TTL = 8 seems to be longer in comparison to the other estimation results obtained. This is probably due to temporary limitations of the ICMP capabilities of router R_8.

Fig. 8 shows a schematic view of the obtained available bandwidth estimates between the endpoint R_0 and routers R_2, R_3, R_6, R_8 and R_9 as well as between the two endpoints R_0 and R_{11}. The available bandwidth estimates used are obtained as average values over the interval 50 – 100 in Fig. 7.

Now, a characterization of the subpaths can be made. The available bandwidth estimate of the path $P(R_0, R_2)$ is 45 Mbps as seen in Fig. 8. This estimate is a lower bound of the available bandwidth of that path, since there may be limitations on ICMP generation at router R_2, as discussed in Section VI.B. Higher TTL values provide further information.

The available bandwidth estimate for the path $P(R_0, R_3)$ is 21 Mbps. Further, the estimate between R_0 and R_6 is 77 Mbps. This indicates that both routers R_2 and R_3 have limitations on ICMP generation, since a subpath cannot have smaller available bandwidth than its superpath. Hence, all subpaths of $P(R_0,,R_6)$ also have an available bandwidth lower bound of 77 Mbps.

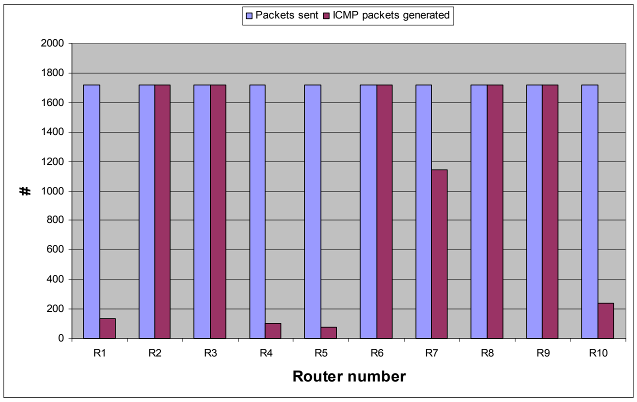

Figure 6. A comparison of the number of ICMP time exceeded packets received and the number of probe packets discarded per TTL value for each router on the Ericsson Research – Uppsala University path.

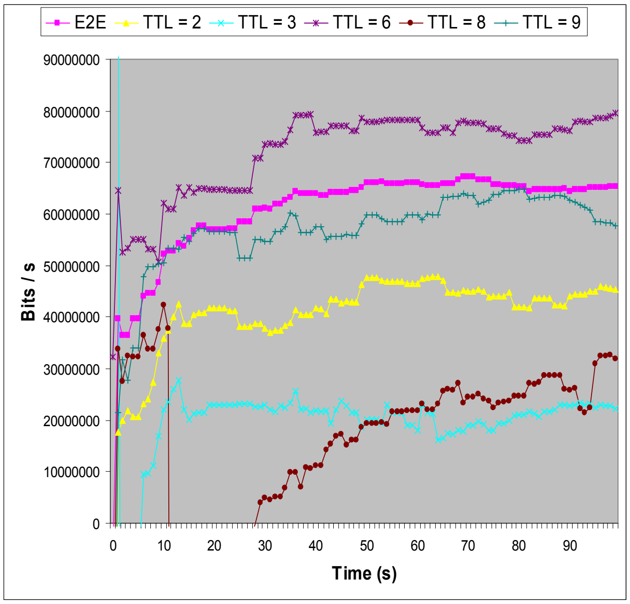

Figure 7. The end-to-end available bandwidth in comparison to available bandwidth estimates for different TTL values over time.

The estimate from R_0 to R_8 is 25 Mbps, and to router R_9 it is 61 Mbps according to Fig. 8. The estimate obtained using R_8 is obviously limited by ICMP generation since a higher available bandwidth estimate is obtained utilizing router R_9. A lower bound of the available bandwidth from the endpoint R_0 to both router R_8 and R_9 is then 61 Mbps.

The end-to-end available bandwidth between R_0 and R_{11} is 67 Mbps as can be seen in Fig. 8. Note that this is not a lower bound, but an estimate of the actual value. Thus, the lower bound of the available bandwidth from the endpoint R_0 to both router R_8 and R_9 may be further sharpened to 67 Mbps. Since a lower bound of the available bandwidth between R_0 and R_6 is 77 Mbps, and the actual value for the whole path is 67 Mbps, it can be inferred that the bottleneck must be located between R_6 and R_{11}.

This successive improvement (sharpening) of the available bandwidth lower bounds for the ICMP-probed subpaths is summarized in Table 1. The left data column contains the raw estimates (lower bounds) of the available bandwidth of each subpath directly probed using the ICMP scheme, as displayed in Fig. 8.

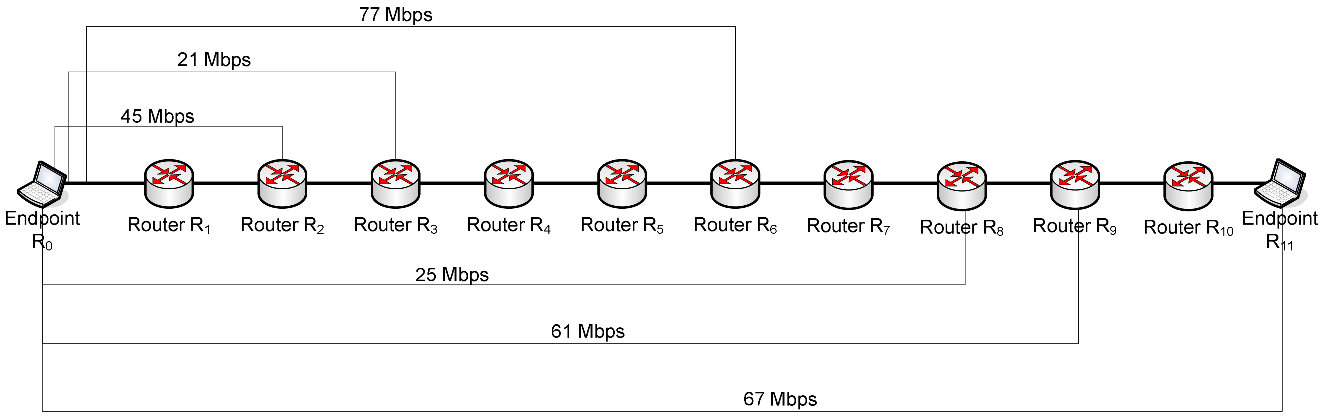

Figure 8. Illustration of the estimates of the available bandwidth for TTL = 2, 3, 6, 8 and 9 on the path between an Ericsson Research node R_0 and a node at Uppsala University R_{11}. The end-to-end available bandwidth is also shown. The estimates are average values over the interval 50 – 100 from Fig. 7.

The second data column shows the lower bounds improved by taking into account the fact that a subpath cannot have a lower available bandwidth than its superpath, i.e. that the sequence of available bandwidths of successive superpaths $\{A(P(R_0,R_k))\}_k$ has to be non-increasing.

Finally, the third data column shows the further improved bounds when also taking the end-to-end measurement into account.

From these data for overlapping paths, it is possible to infer lower bounds for the available bandwidth of the disjoint subpaths $P(R_0,R_2)$, $P(R_2,R_3)$, … , $P(R_9,R_{11})$ that the end-to-end path $P(R_0,R_{11})$ is partitioned into by the ICMP-enabled routers. This is summarized in Table 2.

This partitioning of the end-to-end path may be continued all the way to the individual router hops. For instance, since $P(R_0,R_1)$ and $P(R_1,R_2)$ are subpaths of $P(R_0,R_2)$, and we have the available bandwidth lower bound 77 Mbps for $P(R_0,R_2)$, we know that this lower bound also holds for each of the two subpaths, as neither can have a value lower than that of their superpath $P(R_0,R_2)$. So, we arrive at Table 3 specifying lower bounds for the available bandwidth of each router hop.

TABLE I. LOWER BOUNDS FOR THE AVAILABLE BANDWIDTH OF THE MEASURED SUBPATHS, IMPROVED FIRST BY TAKING THE NON-INCREASING NATURE INTO ACCOUNT FOR THE ICMP-GENERATED VALUES, AND FURTHER BY ALSO HAVING ACCESS TO THE END-TO-END AVAILABLE BANDWIDTH ESTIMATE. NOTE THAT THIS LATTER VALUE (INDICATED BY AN ASTERISK) IS NOT A LOWER BOUND, BUT AN ESTIMATE OF THE TRUE VALUE OF THE AVAILABLE BANDWIDTH.

Subpath	Raw bound (Mbps)	Improved bound (Mbps)	Further improved bound (Mbps)
$P(R_0,R_2)$	45	77	77
$P(R_0,R_3)$	21	77	77
$P(R_0,R_6)$	77	77	77
$P(R_0,R_8)$	25	61	67
$P(R_0,R_9)$	61	61	67
$P(R_0,R_{11})$			67*

TABLE II. LOWER BOUNDS FOR THE AVAILABLE BANDWIDTH OF THE DISJOINT SUBPATHS $P(R_0,R_2)$, $P(R_2,R_3)$, … , $P(R_9,R_{11})$ THAT THE END-TO-END PATH $P(R_0,R_{11})$ IS PARTITIONED INTO BY THE ICMP-ENABLED ROUTERS

Subpath	Bound (Mbps)
$P(R_0,R_2)$	77
$P(R_2,R_3)$	77
$P(R_3,R_6)$	77
$P(R_6,R_8)$	67
$P(R_8,R_9)$	67
$P(R_9,R_{11})$	67

TABLE III. LOWER BOUNDS FOR THE AVAILABLE BANDWIDTH OF EACH ROUTER HOP OF THE END-TO-END PATH $P(R_0,R_{11})$.

Subpath	Bound (Mbps)
$P(R_0,R_1)$	77
$P(R_1,R_2)$	77
$P(R_2,R_3)$	77
$P(R_3,R_4)$	77
$P(R_4,R_5)$	77
$P(R_5,R_6)$	77
$P(R_6,R_7)$	67
$P(R_7,R_8)$	67
$P(R_8,R_9)$	67
$P(R_9,R_{10})$	67
$P(R_{10},R_{11})$	67

VIII. SUMMARY AND CONCLUSIONS

This paper presents a novel probing scheme which can be used for estimating the available bandwidth of subpaths, without the requirement of control over two endpoints of a network path. Instead of a probe-packet receiver, this scheme utilizes the ICMP capability of routers in order to return a stream of ICMP time exceeded packets, corresponding to the stream of probe packets transmitted. By analyzing the time stamps of the transmitted and received packets, an estimate of the available bandwidth from the endpoint to a router is obtained in much the same way as for end-to-end probing methods.

By identifying several ICMP-enabled routers on the path of interest, it is shown to be possible to probe the subpaths from a sending endpoint to each of these routers by choosing the TTL value to correspond to the remote router.

An important observation is that any ICMP packet generation delays will contribute to the time dispersion of the received ICMP packets, which implies that the available bandwidth estimate obtained should be interpreted as a lower bound.

These lower bounds for overlapping subpaths may be sharpened, by taking into account that the sequence of available bandwidths of successive superpaths has to be non-increasing.

Further, if one has access to a remote endpoint and can perform end-to-end measurements, it may be possible to further sharpen the bounds. Also, it may be possible to identify a portion of the end-to-end path which contains the bottleneck.

The set of ICMP-enabled routers on the path from a sending endpoint to a remote network node defines a partition of that path into disjoint subpaths. By taking into account the fact that the available bandwidth of a superpath cannot be larger than that of a subpath, it is possible to infer lower bounds for the available bandwidth of these disjoint subpaths.

Finally, these can be further decomposed into the individual router hops, and a lower bound can be assigned to each.

In order to demonstrate the feasibility of this ICMP-based scheme, the paper reports on a measurement study conducted over the Internet. First a set of Swedish routers was investigated to study whether a sufficient number of them generate the required ICMP packets. It was seen that about half of the routers actually respond to all “requests” and thus reply with ICMP time exceeded packets for each discarded probe packet. This indicates that except for the shortest paths, typically it is possible to identify a set of ICMP-enabled routers on the path, and thus to employ this scheme.

Secondly, one Swedish Internet path between Ericsson Research in Stockholm and Uppsala University was investigated in more detail in order to illustrate this technique.

It is concluded that the ideas presented in this paper could lead to additional tools to be used for characterising network paths in greater detail, in terms of available bandwidth, than is possible using pure end-to-end methods.

REFERENCES

[1] Melander, B., Björkman, M., Gunningberg, P.: Regression-based availale bandwidth measurements. In: Proceedings of the 2002 International Symposium on Performance Evaluation of Computer and Telecommunications Systems, San Diego, CA, USA. (2002)

[2] Liu, X., Ravindran, K., Loguinov, D.: Multi-hop probing asymptotics in available bandwidth estimation: Stochastic analysis. In: Proceedings of ACM SIGCOMM Internet Measurement Conference. (2005)

[3] Prasad, R., Murray, M., Dovrolis, C., Claffy, K.: Bandwidth estimation: metrics, measurement techniques, and tools. In IEEE Network. (2003)

[4] IETF RFC 5136: Defining Network Capacity. (2008)

[5] ITU-T Y.1540 Amendment 1: IP-layer Capacity Framework. (2009)

[6] Lee, Y., Agarwal, S., Butcher, C., Padhye, J.: Measurement and estimation of network qos among peer xbox 360 game players. In: Passive and Active Measurement Conference. (2008)

[7] Ekelin, S., Nilsson, M., Hartikainen, E., Johnsson, A., Mångs, J., Melander, B., Björkman, M.: Real-time measurement of end-to-end available bandwidth using kalman filtering. In: Proceedings to the IEEE/IFIP Network Operations and Management Symposium, Vancouver, Canada. (2006)

[8] Ribeiro, V., Riedi, R., Baraniuk, R., Navratil, J., Cottrel, L.: pathchirp: Efficient available bandwidth estimation for network paths. In: Passive and Active Measurement Workshop. (2003)

[9] Harfoush, K., Bestavros, A., Byers, J.: Measuring bottleneck bandwidth of targeted path segments. In: Proceedings of Infocom: The IEEE International Conference on Computer Communication, San Fransisco, USA. (2003)

[10] Mah, B.: pchar (2005) http://www.kitchenlab.org/www/bmah/Software/pchar/

[11] Dovrolis, C.; Ramanathan, P.; Moore, D.: Packet-dispersion techniques and a capacity-estimation methodology. In: IEEE/ACM Transactions on Networking, Volyme 12 (2004)

[12] IETF RFC 792: Internet Control Message Protocol, DARPA Internet Program Protocol Specification (1981)