On Measurement of the End-to-End Available Bandwidth Across t he Internet

Reza Rejaie
Information and Computer science Department
University of Oregon
reza@cs.uoregon.edu

I. Introduction

Available bandwidth is one of the key properties of an end-to-end path across the Internet. Although the term "available bandwidth" is frequently used in the networking literature, it still does not have a clear and widely accepted definition and therefore it is not an easy property to measure. To clarify this issues, a recent work by Jain et al [1] have defined available bandwidth as an unused bandwidth at the bottleneck link along an end-to-end path and presented a mechanism to measure this unused bandwidth. However, as we will discuss in this paper, such a definition is not useful.

In this position paper, we argue that *the available bandwidth between two end points across the Internet is the fair share of the bottleneck link's bandwidth that is estimated by a network-friendly congestion control (CC) mechanism*. For the rest of this paper, first we argue why this definition of available bandwidth is the most reasonable and useful one for the Internet. Second, we describe how this definition can be employed to estimate available bandwidth across a path and how this information should be used by an application. Third, we identify a couple of key issues and challenges in estimating available bandwidth.

II. How to Define Available Bandwidth

The basic definition of the available bandwidth for an Internet path is the bandwidth that can be obtained by a transport protocol across that path. The available bandwidth to each flow primarily depends on the share of that flow's bandwidth at a bottleneck link which is the slowest link across the end-to-end path. Since Internet resources (*i.e.,* bandwidth) are shared, all flows are expected to perform congestion control. Performing congestion control by each flow achieves in three important goals: *(i)* it ensures stability of the network, *(ii)* coexisting flows obtain a fair share of the bottleneck link bandwidth, and *(iii)* it ensures high utilization of network bandwidth. Ideally, a single congestion controlled flow should be able to fill up the pipe and fully utilize bandwidth of a bottleneck link. As the number of co-existing flows at the bottleneck increases, their CC mechanisms compete for the bottleneck bandwidth until their reach a new equilibrium where each flow obtains an even (or RTT proportional) share of the bottleneck bandwidth. In practice, utilization of a bottleneck link and inter-protocol fairness depend on several factors such as queue management and the buffer size at the bottleneck, level of statistical multiplexing (*i.e.,* number of coexisting flows) as well as deployed CC mechanisms by coexisting flows. For example, tens of TCP flows (with sufficiently large upper bound for congestion window so that flow control does not limit transmission rate) can easily fill up a single bottleneck link with sufficient buffer size (*e.g.,* 3*RTT*BW). This high level description of interactions among flows in the Internet leads to two key points: First, in an environment such as the Internet where a majority of the flows (at least TCP traffic) should be congestion controlled, most of the bottleneck links are likely to be fully utilized. The obtained bandwidth by a new flow f depends on the interaction between the congestion control mechanism deployed by f and other coexisting flow. In theory, no matter how congested a bottleneck link might be, a new congestion controlled flow can obtain a portion of the bottleneck bandwidth. This implies that the defining available bandwidth as an unused capacity of a bottleneck link is not appropriate since such unused bandwidth may not exist even when the available bandwidth to a new flow is not zero. Clearly, if such an un-utilized bandwidth exists on a given link, that link is not considered a bottleneck. Second, a network-friendly congestion control mechanism should be designed such that it achieves both inter- and intra-protocol fairness under variety of network conditions (*e.g.,* different level of multiplexing, mix of traffic), *i.e.,* the CC mechanism should achieve neither more nor less than its fair share of bottleneck bandwidth. This is indeed the key evaluation of any new CC mechanism. Therefore, such a CC mechanism is not intrusive and it can effectively estimate a fair share of bottleneck bandwidth which is considered "available bandwidth" to this flow. This is in fact the actual bandwidth that is obtained by this flow.

III. How to Measure & Use Available Bandwidth

Conceptually, a CC mechanism continuously probes the path to estimate the available bandwidth. Then, the available bandwidth should be reported to the application in order to match its transmission rate accordingly. A key question is "what is the proper timescale for estimating and reporting the available bandwidth?" The congestion controlled bandwidth depends on offered load by the coexisting flows at the bottleneck link. This implies that the available bandwidth can exhibit unpredictable and potentially wide variations over short timescales (*e.g.,* once per RTT). Therefore, the application should match its transmission rate over these short timescale, otherwise it may interfere with the behavior of the CC mechanism. To illustrate this point, consider an example where a CC mechanism reports the available bandwidth once per RTT but the application matches its transmission rate with average available bandwidth over N RTT. This means that the actual transmission rate might be higher than the available bandwidth in some RTTs and less than the available bandwidth in other RTTs. In this example, the application can significantly change pattern of packet transmission during N*RTT period. This could in turn affect accuracy of the estimated available bandwidth by the CC mechanism since the transmitted data packets are used by the CC mechanism to probe. An extreme example is an application that sends a 10*RTT worth of available bandwidth in one or multiple bursts. We conclude that *(i)* the congestion controlled bandwidth should be reported and matched by the application over short timescale (*e.g.*, a small number of RTT), and *(ii)* the application should minimize the variations in transmission pattern to have minimal impact on the accuracy of CC mechanism. We note that the extent of any potential impact by application on the accuracy of the CC mechanism has not been sufficiently studied and requires further investigation.

IV. Remaining Challenges

Defining available bandwidth as the estimated bandwidth by a CC mechanism tightly couples behavior of the available bandwidth (specially over short timescales) to the dynamics of the employed CC mechanism. For example, TFRC[2] exhibits smoother variations in available bandwidth in compare with TCP or RAP [3]. This implies that the available bandwidth can be estimated in different ways as long as the overall behavior remains network friendly.

One interesting issue that deserves further investigation is the statistical properties of congestion controlled bandwidth across the Internet. For example, for a given end-to-end connection, it would be interesting to answer the following questions: "what are typical variations in available bandwidth?", "given some definition of congestion, how often congestion occurs?", "how long does a congestion event typically last?", "Do these characteristics depend on the level of multiplexing across the bottleneck link?", "Can these characteristics be captured by a model?", "Do we observe statistically similar behavior across different end-to-end paths?".

References

[1] M. Jain and C. Dovrolis, “End-to-end available bandwidth: Measurement metodology, dynamics, and relationship with tcp throughput,” in *ACM SIGCOMM*, 2002, reviewed.

[2] S. Floyd, M. Handley, J. Padhye, and J. Widmer, “Equation-based congestion control for unicaqt applications,” in *Proceedings of the ACM SIGCOMM*, 2000.

[3] R. Rejaie, M. Handley, and D. Estrin, “RAP: An end-to-end rate-based congestion control mechanism for realtime streams in the internet,” in *Proc. IEEE Infocom*, New York, NY., Mar. 1999.