Admission Control for IP Quality of Service

Yaw Opoku, JinSoo Park, Todd J. Eshler, Kevin P. Flanagan
Bradley Department of Electrical and Computer Engineering

Abstract

This paper introduces different types of admission control for IP Quality of Service. Three different approaches for admission control are introduced: measurement-based admission control, endpoint admission control, and policy-based admission control. The benefits of each approach are discussed relative to quality of service in an IP network. Some potential applications are also discussed. An example protocol used for RSVP-based admission control is then described. This protocol demonstrates a mapping from the internet level to an IEEE 802-style LAN.

Introduction

In the past several years, the Internet Engineering Task Force (IETF) working groups have been trying to extend the IP protocol to support Quality of Service (QoS) better. IntServ and RSVP have been developed to help provide end-to-end QoS. A simpler implementation, DiffServ, has come about that has more scalability than IntServ. Both of these implementations and many others need a process to determine if a traffic flow should be admitted or rejected. This process is called admission control. Admission control can be based on network measurements or network policies and can manage flows or groups of flows. Many different techniques and protocols have been developed to support admission control for IP networks.

1. Measurement-based Admission Control [Yaw Opoku]

Introduction

One of the biggest challenges of Admission Control for IP QoS is sharing network resources efficiently among delay-sensitive flows that are admitted into the network. Traditional Admission Control Algorithms (also known as parameter-based admission control) mainly rely on flow-specified statistical models (e.g., M/M/1, M/D/1, M/M/n, etc.) that are used to obtain parameters that describe the traffic characteristics of the flows. The parameters are then used by the flows to estimate the amount of resources they need. The network then uses these estimates to admit or reject flows depending on availability of the requested resources.

But the problem with this approach is that it is often very difficult for applications to come up with models that accurately describe their traffic patterns. This is even more challenging in an IP network since most of the traffic in such networks are significantly heterogeneous. Therefore in order to obtain guaranteed resources (when admitted), flows provide *a priori* parameter estimates that represent their worst-case behaviors. This causes over-booking of network resources, which results in low network utilization. Utilization is even lower when the flows are bursty.

Measurement-based Admission Control

Measurement based admission control (MBAC) is an alternative to parameter-based admission control, whose admission decisions are mainly based on traffic parameter estimates from measurements obtained from an existing traffic. Every MBAC has two components, namely:

1. A measurement procedure for estimating current network load, and
2. An algorithm that uses the load estimate to make admission decisions. In the following sections, I will describe two admission algorithms and two measurement procedures that illustrate the MBAC concept fairly clearly:

Admission Algorithms

1. Measured Sum:

Let be μ be the link bandwidth, r^{α} the amount of bandwidth requested by flow α, and $\hat{u}$ be the measured load of current traffic. Then the measured sum algorithm accepts α into the network only if the following condition is satisfied:

$$\hat{u} + r^{\alpha} < c\,\mu$$

where $0 < c < 1$; c is a user-specified utilization target. When the bandwidth utilization approaches 100%, the variation of packet delay grows large and the algorithm will make wrong admission decisions as a result. Therefore a utilization target is set and the network is kept operating under the desired utilization.

2. Hoeffding Bounds:

This algorithm calculates the equivalent bandwidth C, based on Hoeffding bounds [7], of the flows on a link. In reference [10], $C(\varepsilon)$ is defined as the probability that the steady-state bandwidth requirement of n flows exceeds ε, assuming the flows have been policed to have a peak rate of p. C is given by:

$$\hat{C}_H(\hat{v}, \{p_i\}, \varepsilon) = \hat{v} + \sqrt{\tfrac{1}{2}(\ln(1/\varepsilon)\textstyle\sum_i (p_i)^2}$$

where $\hat{v}$ is the measured average arrival rate of current traffic. The flow α is admitted if the following condition is satisfied:

$$\hat{C}_H + p^a \le m$$

where p^a is the peak rate of α, and μ is the link bandwidth. The current average arrival rate is then updated as follows: $\hat{v}' = \hat{v} + p^a$. It should be noted that the above algorithms assume the existence of some signaling protocols (e.g., RSVP) that are used by flows to request QoS.

Measurement Procedures

1. Time-windows:

This measurement procedure divides time into windows of size T, and within each window, average load is measured with a sampling period of S. The highest load estimate of a window is used as the load of the next window. If a new estimate is found to be greater than the current load, the load is immediately replaced by the new estimate.

2. Exponential Averaging

In this procedure, the packet arrival rate, R, of a flow is measured with a sampling period of S. The new average arrival rate of the flow is then computed with the infinite impulse response function:

$$\text{AvgR}' = (1 - w)*\text{AvgR} + w*R$$

where w is a user-specified weight.

Conclusion

Several MBACs have been proposed in recent years due to their promise to offer better support for real-time applications, and simulation experiments of these algorithms have indicated that they achieve higher utilization than parameter-based admissions control for a given QoS. However, all the MBACs that have been proposed seem to yield very similar link utilizations, and this has raised questions about the direction of future research in MBAC [4].

2. Endpoint Admission Control [JinSoo Park]

Introduction

As an alternative to the IntServ algorithm, endpoint admission control has been introduced. The IntServ achieves individual QoS in IP network on per-flow bases by using a RSVP as means to reserve resources in the network from source to destination. However, it has a scalability problem, since routers need to retain state information and reserve resources along the way. Meanwhile, endpoint admission control algorithm does not depend on the routers for the admission control. Therefore, routers do not need to keep per-flow state or process reservation request and can drop or mark packets for some other QoS related-purposes.

Issues Regarding the Endpoint Admission Control Algorithm

- What type of queuing algorithm is better?

Fair queuing or its variants are not appropriate to service the admission-controlled traffic, since their flow acceptance could impair the service being delivered to the already accepted larger flows.

- What should be done when the admission-controlled flows coexist with best-effort traffic?

In this case, we need an algorithm that does not allow the admission-controlled traffic to borrow bandwidth from best-effort, and does not allow the best-effort traffic to pre-empt admission-controlled traffic. We can achieve this by placing admission-controlled traffic in a higher priority class and strictly limiting its share of bandwidth to some fraction of the link bandwidth. This could be implemented as a simple priority queue with a rate limiter in routers. However, such queuing mechanisms may leave the link idle temporarily rather than send the admission-controlled traffic even when there is no best-effort traffic present.

- How can multiple levels of service be achieved under priority queuing mechanisms?

We can achieve multiple levels of admission-controlled service only as long as the probe packet is placed at the lower priority than all other admission-controlled traffic so that we can send the data packets at the different levels of priority. It can be simply done by using a different DS field for the probe packets than for the data packets.

- How can the acceptance threshold be selected?

Endpoint admission control will show best performance if all flows adopt the same acceptance threshold, because the quality of service of all flows depends on the least stringent acceptance threshold of any flow. Therefore, flows have little to gain by choosing a more stringent acceptance threshold.

- What is the in-band and out-of-band probing?

In-band probing sends probe packets together with the data packets, but we may need to choose an option between long set-up times and small loss fractions, or short set-up times but somewhat higher loss fractions. Out-of-band probing sends probe packets at a lower priority than the admission-controlled traffic, and can achieve a suitable size of acceptance threshold with a reasonable set-up delay and lower data losses.

Probing Types

- The Simplest Form of Probing

The sending and receiving hosts work together; the receiving host calculates the losses and communicates the acceptance/rejection decision to the sending host. If the number of packet losses exceeds the thresholds even before the end of probing, then probing stops and the flow will be rejected.

- The Slow-start Probing

When the input load is much higher than the link capacity, the system may experience a collapse in in-band probing, or a starvation in out-of-band probing. Therefore it may need to ramp up the probing rate slowly in order to prevent these events.

- The Early Rejection Probing

If the loss percentage goes above limits in any second-long interval, then the flow will be rejected in order to determine the performance differences due to early rejection or the incremental increases.

Summary and future work

The endpoint admission control algorithm was contrived to overcome the scalability problem of the IntServ. The host (end point) decides whether or not to admit the flow based on the calculated loss percentage of probe packets. It does not require any router support; therefore, the routers do not need to keep state information. However it has some substantial problems, such as long set-up delays that may limit its appeal for certain applications. Also utilization and loss percentage degrade in a substantial degree under sufficiently heavy loads, even with slow-start probing. More importantly, design and deployment challenges still lie ahead. Further research in this arena needs to be performed, but the stakes seem to be very high.

3. Policy-based Admission Control [Todd J. Eshler]

Traditionally, the IP architecture provides best-effort services to connections in a network. RSVP, IntServ, and DiffServ have been developed by IETF working groups to add QoS extensions to provide other services besides best-effort. All of these extensions require some form of admission control. Usually, admission control determined admission of flows by taking into account the amount of network resources available versus the amount of network resources currently in use. The idea of policy-based admission control [12] has risen. Policy-based admission control not only includes traditional admission control but also takes into account network policy criteria. These criteria include policy based on time that the request is made, user or application making the request, or other security considerations.

The IETF Network Working Group has provided a framework for policy-based admission control [12]. In this framework, requirements for a policy-based admission control [12] have been setup. These requirements are:

- Preemption: the protocol should allow for an old request to be replaced by a new request.
- Multiple styles of policies: The protocol should allow for more than one style of policy to be carried out. This allows for network administrators or service providers to determine what policies should be enforced on a network.
- Fault Recovery: The protocol should be able to help enforcing policies if a node or link fails. Also the protocol should be able to handle structural changes in the network.
- Policy Ignorant Nodes (PINs): The protocol should still be able to enforce policies on a network if there are nodes that do not support the protocol in the network. This allows for legacy hardware to still be used in a network.
- Scalability: The protocol should be able to handle multiple flows and multiple nodes.
- Security: The protocol should provide some form of security to prevent denial of service attacks and theft of services.
- Monitor Account Information: The protocol should be keep track of account information such as policy state, resource usage. The account information provided by the protocol could be used for billing services.

Common Open Policy Service Protocol (COPS) [1] is an example of a policy-based admission control protocol that has been developed by the RSVP working group to provide policy-based admission control in networks using RSVP [1] [2] [9]. The COPS protocol runs on top of TCP to distribute objects [1]. These objects contain policy and admission control information. Context objects are used to keep a policy decision point (PDP) informed of RSVP requests and admissions made by a policy enforcement point (PEP) [2].

As IP networks evolve to support more QoS features, policy-based admission control will play a large role in what services users will have.

4. Subnet Bandwidth Manager [Kevin P. Flanagan]

Extensions to the Internet architecture have been defined so that applications can request specific qualities of service from an internetwork. This allows an improvement to the current Internet Protocol (IP) best-effort service. Examples of these extensions include Resource Reservation Protocol (RSVP) and definition of new service classes to be supported by Integrated Services (IntServ) routers. Since these extensions are independent of the underlying networking technologies, "it is necessary to define the mapping of RSVP and Integrated Services specifications onto specific subnetwork technologies" [11]. This section will present Subnet Bandwidth Manager (SBM), a signaling protocol for RSVP-based admission control over IEEE 802-style networks. SBM provides a method for mapping RSVP onto IEEE 802-style networks [11].

Though the current trend is toward a greater usage of switched Local Area Network (LAN) topologies that support priority queuing specified by IEEE 802.1p, it should be safe to assume that LAN technologies will continue to be a mix of legacy shared/switched LAN segments and newer switched segments. The SBM protocol would allow for a range of solutions that vary from "purely administrative control (over the amount of bandwidth consumed by RSVP-enabled traffic flows) to one that requires cooperation (and enforcement) from all the end-systems in an IEEE 802 LAN" [11].

The main procedure for SBM-based admission control requires first an entity called the Designated SBM (DSBM). The presence of a DSBM makes a segment "managed". The DSBM is responsible for admission control over the reservation requests from the DSBM clients in that segment. Multiple SBMs may exist on a given segment, but a single DSBM is elected for the segment.

Once a DSBM is selected, it is initially configured with information such as the limits on fraction of available resources that can be reserved on each managed segment. The configuration is likely to be static with current devices, but future work may allow for dynamic discovery of this information. For each attached interface, a DSBM client determines whether a DSBM exists on the interface. This is necessary to determine is the client can communicate with the DSBM for admission control purposes [11]

To request reservation of resources, a DSBM client sends or forwards an RSVP PATH message over a managed segment. Instead of sending to PATH message to the destination address, the client sends it to the DSBM. After processing, the DSBM will forward the PATH message toward its destination. Later, when the RESV message is sent back along the path, it is sent to the previous message stored in the PHOP object, now the DSBM address. In this manner, the DSBM can control all reservations going through its managed segment(s).

One enhancement SBM has to traditional RSVP signaling is the storage of both Layer-2 (L2) and Layer-3 (L3) addresses for the HOP objects. In the case where a switch must forward a request, the switch does not have the capability to determine the MAC address based solely on the IP address stored in the LAN_NHOP object. Therefore, two new objects have been created to store this necessary information: LAN_NHOP_L2 and LAN_NHOP_L3.

Another important enhancement is the ability to eliminate duplicate PATH messages when the RSVP session address is a multicast address. Using the LAN_LOOPBACK object, DSBM clients place their own unicast IP address in the object, so that when receiving another PATH message it can determine if a loop has occurred.

SBM has several interoperability concerns. An L2 domain connected to the network may not have RSVP capabilities. There could potentially be devices that are not SBM-compatible connecting two segments. Some RSVP senders may not be DSBM clients. Finally, a non-SBM router may connect two L2 LANs. In order for SBM to be a reasonable solution for admission control, all of these issues must be addressed.

In the case where a particular domain does not use RSVP, the SBM protocol may co-exist, but will not attempt any admission control. If a device that is not SBM-compatible connects two managed segments, it must be treated as one managed segment. Only one DSBM will be chosen to manage admission control over both segments. This is critical for older-style switches. Potentially, there could be RSVP senders that are not DSBM clients. This will not interfere with the protocol as long as these senders are “DSBM aware” [11]. In other words, these senders must know that SBM is being used, and be aware that the DSBM will insert its own IP address in the path for RSVP messages and responses. Finally, when non-SBM routers connect L2 domains, SBM messages must not pass from one domain to the other. If all of the interoperability concerns are handled, SBM has the benefit of operating among many types of routers and switches, thus increasing the flexibility of the protocol. This allows SBM to vary from total admission control to potentially being transparent for devices, which are not capable of supporting such a protocol.

Conclusions

The best method or approach to admission control for IP QoS is still in debate. Some of the best-known methods and protocols have been shown, and techniques for analyzing them are given. The best solution is a blend of the types of admission control described in this report. Future work is needed to standardize the approach to allow the greatest amount of control with flexibility to work with other protocols and implementations for network QoS.

References

[1] Boyle, J., Cohen, R., Durham, D., Herzog, S., Raja, R. and A. Sastry, "The COPS (Common Open Policy Service) Protocol", RFC 2748, January 2000.

[2] Boyle, J., Cohen, R., Durham, D., Herzog, S., Raja, R. and A. Sastry, "COPS Usage for RSVP", RFC 2749, January 2000.

[3] Breslau, L.; Jamin, S.; Shenker, S. "Comments on the Performance of Measurement-based Admission Control Algorithms". INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE , Volume: 3 , 2000 Page(s): 1233 -1242 vol.3.

[4] Breslau, L.; Jamin, S.; Shenker, S. "Measurement-based Admission Control: What Is the Research Agenda ?" Quality of Service, 1999. IWQoS '99. 1999 Seventh International Workshop, 1999 Page(s): 3 –5.

[5] Breslau, L., Knightly, E.W., Shenker, S., Stoica, I., and H. Zhang, "Endpoint admission control: Architectural Issues and performance", *in Proceedings of Sigcomm 2000*, (Stockholm, Sweden), Aug.2000.

[6] Cetinkaya, C. and E. Knightly,"Egress Admission Control," *Proc. IEEE INFOCOM*, March 2000.

[7] Floyd, S., "Comments on Measurement-based Admissions Control for Controlled Load Services". Technical report, Lawrence Berkeley Laboratory, July 1996.

[8] Grossglauser, M. and David N.C. Tse, "A Frame Work for Robust Measurement-based Admission Control", IEEE/ACM Trans. On Networking, Vol. 7, No.3, June 1999.

[9] Herzog, S., "RSVP Extensions for Policy Control", RFC 2750, January 2000.

[10] Jamin, S., Shenker, S.J., Danzig, P.B. "Comparison of Measurement-based Admission Control Algorithms for Controlled-load Service". INFOCOM '97. Sixteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Driving the Information Revolution, Proceedings IEEE, Volume: 3, 1997 Page(s): 973 -980 vol.3

[11] Yavatkar, R., et al., "SBM (Subnet Bandwidth Manager): A Protocol for RSVP-based Admission Control over IEEE 802-style networks", RFC 2814, May 2000.

[12] Yavatkar, R., Pendarakis, D. and R. Guerin, "A Framework for Policy-Based Admission Control", RFC 2753, January 2000.