

Improvement of performance in MPLS domains by using Caching and Aggregation of CR-LSP

Franco Tommasi, Simone Molendini and Andrea Tricco
Dept. of Innovation Engineering and ISUFI, University of Lecce, ITALY

Abstract: Constraint-based routing (CR-LDP) extends the features of MPLS Label Distribution Protocol (LDP) working in conjunction with a wide variety of QoS scenarios and services including IP differentiated services, integrated services, traffic engineering, Asynchronous Transfer Mode (ATM) service classes and frame relay. When a large number of CR-LSPs are required a lot of CR-LDP signalling is needed: this paper proposes two techniques - caching and aggregation - to reduce the number of CR-LDP messages by dynamically increasing the number of flows served by each single CR-LSP. The results of a simulation are also used to discuss the performance of such mechanisms.

Key words: QoS, CR-LDP, scalability.

1. INTRODUCTION

The MultiProtocol Label Switching (MPLS) protocol allows high-performing label switching of packets [1][2]. Network traffic is forwarded into the domain using a simple label: into an MPLS domain two Label Switching Routers (LSRs) must agree on the meaning of the labels used by the forwarded traffic. The Label Distribution Protocol (LDP) distributes labels between LSRs to allow MPLS forwarding along normally routed paths [3]. Constraint-based routing (CR-LDP) extends the features of the LDP protocol [4] working in conjunction with a wide variety of Quality of Service (QoS) scenarios and services including IP differentiated services, integrated services, traffic engineering, Asynchronous Transfer Mode (ATM) service classes and frame relay [5][6].

The creation and release of a constraint-based routed label switched path (CR-LSP) is based upon three CR-LDP messages: Label Request, Label

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: [10.1007/978-0-387-35620-4_43](https://doi.org/10.1007/978-0-387-35620-4_43)

Mapping and Label Release. The Ingress LSR must send a Label Request towards the Egress LSR to ask each node along the path to create the desired CR-LSP. This message contains the Traffic Parameters TLV (Type Length Value) specifying the QoS parameters associated with the path. After having received this message, the Egress LSR sends a Label Mapping message in the opposite direction to confirm the creation of the path in each LSR. A CR-LSP may be released by the Label Release message. When a large number of CR-LSPs are required a lot of CR-LDP signalling is needed [7].

The purpose of this paper is to propose some solutions that improve the scalability of the CR-LDP protocol and reduce the size of label's state information a router must hold. Such solutions are dynamically applied within a given Ingress/Egress couple to diminish signalling overhead. The rest of the paper is organized as follows. Section 2 and 3 introduces the use of caching and aggregation with CR-LSP. Section 4 discusses a dynamic mechanism for using such techniques. Final conclusions are provided in Section 5.

2. HOW TO CACHE CR-LSPS

No modification to the current specifications of the CR-LDP protocol is required in order to benefit the MPLS LSRs using the caching we describe in this paragraph. Only the Ingress LSR at the border of the MPLS domain must be aware of the use of this technique.

Let the name of a flow be F1. According to MPLS architecture, Ingress LSR is the only LSR responsible for sending creating and releasing CR-LSP and it is also the only LSR responsible for the association between flows' packets and the labels' chain used along that path. When a request to release F1 arrives to the Ingress LSR, this node may decide not to transmit the Label Release message into the MPLS domain. The Ingress LSR waits for another flow, let's say F2, directed to the same Egress LSR and with similar QoS, for a maximum given amount of "C" minutes. To forward the traffic of the new flow (F2) the Ingress LSR only changes the Forwarding Equivalence Class (FEC) associated to the CR-LSP. Instead of sending a Label Release message for F1 and a Label Request message for F2, the Ingress LSR transmits no message into the MPLS domain and the LSRs along the path hold that CR-LSP. This mechanism may be arbitrarily generalised to N flows using, at different times, the same CR-LSP without signalling. An exception is made for the original Label Request message. The CR-LSP is definitely released from the domain after an idle period of C minutes. This mechanism looks very much like a caching strategy.

The cache allows different flows to use, at different times, the same CR-LSP: what about the flows' QoS requirements? F2 will receive inside the domain a service similar to the expected one only if its Traffic Parameters TLV are similar to those of F1. Therefore, we introduce a parameter "A" that is the degree of affinity between the traffic characteristics associated to the CR-LSP and the parameters of the flow wanting to use it.

3. HOW TO AGGREGATE CR-LSPS

The basic idea behind this technique is to allow an Ingress LSR to dynamically aggregate towards the same Egress LSR more than one flow into a single CR-LSP. Thanks to the smaller number of CR-LSPs used, less CR-LDP signalling is needed. To exploit this benefit the CR-LSP's traffic parameters must allow it to carry all the aggregated flows.

The aggregation of flows needs the following steps. At the arrival of a setup request the Ingress LSR sets up a CR-LSP able to contain "B" (B integer number) distinct flows. This way the Admission Control module at the Ingress LSR reserves more resources than effectively used at this time. The purpose of this mechanism is to advance the setup of a later request: the global reservation satisfies B simultaneous reservation slots – each slot being associated to a single flow and B slots composing such simple CR-LSP. The arrival of a setup request may allow a different flow to occupy a slot of this CR-LSP only if this flow has QoS requirements similar to those of the original flow – and to those of the flows in the other slots: no CR-LDP messages are sent across MPLS domain either to notify such setup or to release it. Once all the flows are released, the CR-LSP itself is released sending a regular Label Release message. This mechanism may be optimised by combining it with the caching: after the removal of the last flow the CR-LSP is kept alive for a period of C minutes before its release.

4. THE ABC PARAMETERS

In the previous paragraphs we have introduced three parameters: A, B, C (ABC parameters from now on) that all together control the trade-off between the reduction of CR-LDP signalling and the time resources are held though not used.

The A ("Affinity degree") parameter measures the degree of similarity between a link/path and a flow. Let us consider three scalar values S1, S2 and A where A is between 0 and 1. We define a scalar flow S2 be "a nice

non-degradation” of a path S1 using A if: 1) S2 is less than S1 and 2) S2 is bigger than $A \cdot S1$. Such non-symmetric definition allows us to control the extent to which the use of the path S2 by the flow S1 is considered advantageous. Remembering that a CR-LDP Traffic Parameters TLV is defined by the 5-tuple (PDR, PBS, CDR, CBS, EBS), we define that the flow F2 is a nice non-degradation of the CR-LSP F1 using A if each parameter of F2 is a nice non-degradation of the related parameter in F1.

The B (“Bandwidth multiplier”) parameter is the bandwidth multiplier factor. When an incoming reservation request arrives, the Ingress LSR sets up a CR-LSP able to contain B distinct similar flows as showed in the previous paragraph.

The C (“Caching time”) is the amount of time a CR-LSP should be kept alive after it is no longer used by any flow. This way the Ingress LSR controls the wait time for a new flow.

The choice of the ABC parameters greatly affects the dynamics of the management of the resources. The administrator of the network should set these parameters for all the interfaces of the LSRs in the domain. The values may change for each interface - some may be connected to a well-provisioned network, some might not. To simplify the administration work, a default configuration should be given by the vendor that takes into account the features of the product. We discuss now how the configuration of the ABC parameters modify the trade-off from “conservative” (when less bandwidth is available) to “aggressive” one (when a strong reduction of signalling is desired). Following are some of possible scenarios.

One-to-one direct mapping: $A=(any)$, $B=1$, $C=0$. Neither caching nor aggregation are enabled, hence no A value is needed. This scenario lead to a zero waste of resources but with complete signalling inside the domain. This is the most conservative strategy.

Isolated flows on cached paths: $A=Low$, $B=1$, $C=High$. Caching only is enabled, with a loose affinity value. This way CR-LSPs with large resources may be assigned to flows, that continues to be isolated from each other. The low affinity value grants a large usage of the cached paths. This leads to some wasting of resources, but grants a minimum of signalling in all the cases where reservation and deletion requests incomes with a regular rate.

Aggregation of little flows: $A=Low$, $B=High$, $C=0$. Caching is disabled and aggregation heavily used. This set of values is useful when a number of requests is expected that reserve little resources each. This way, a path is defined for many flows with a single signalling per aggregated path.

Multiplexing on fixed channels: $A=Low$, $B=High$, $C=Infinite$. A set of large (B is High) static channels (infinite caching time) is created with traffic chosen traffic parameters. Incoming traffic is very likely (because of the loose affinity) aggregated to one of them. Otherwise, it will not receive QoS

inside the domain. This configuration needs no signalling inside the domain, except the initial one.

4.1 Simulation

The purpose of this simulation is to study the performances of caching and aggregation mechanisms. We considered QoS flows starting from a single Ingress LSR and directed to the same Egress LSR because the mechanisms introduced in this paper involve a single Ingress/Egress couple. The arrival of QoS requests to the Ingress LSR was modeled as a Poisson process because it simply models a generic arrival of events. The interarrival times follow then an exponential distribution – and we used a parameter λ equal to 1 sec^{-1} . Each setup request contains a single Traffic Parameter which is characterized by a Uniform distribution with a range of [5 kbps, 15 kbps]. The flow's duration time has another exponential distribution using a parameter λ equal to $(180 \text{ sec})^{-1}$. For this study, the A parameter was fixed at the 90% value. We collected the signalling reduction and the waste of bandwidth for the following B and C values: for B, we chose the values {1,2,4,8} and for C the values {0 sec, 100 sec, 200 sec, 400 sec, 600 sec, 800 sec, 1000 sec}.




Figure 1. Signalling/Bandwidth trade-off

Fig. 1 uses the signalling/bandwidth coordinates to show the advantages and final consequences of the methods introduced by this paper. Each point in the plane shows the waste of resources (x axis) and the reduction of signalling (y axis). The lines are plotted for different values of the C

parameter. An increased amount of caching moves the points on the bottom (because it indicates a reduced signalling) and on the right (because it indicates a bigger waste of resources) of the plane. Different lines indicate different aggregation values: a greater number of B slots indicate a more massive use of the aggregation. It is interesting to observe that the tail of the B=1 line is up the head of the B=2 line: this indicates that, while the waste of resources is the same, the line B=2 has a greater reduction of signalling. This picture is an important tool for network administrators to decide the amount of resources that can be dedicated to the reduction of signalling.

5. CONCLUSIONS

A description of some extensions of the CR-LDP protocol is provided in this paper. These mechanisms allow a significant reduction of the number of CR-LDP setup/release messages and, in this sense, they improve the scalability of the CR-LDP QoS architecture. The ABC parameters customize these extensions, and allow tuning the performances with a wide variety of architectures and requirements. The costs of such benefits are an increased complexity of the CR-LDP architecture (but only in the Ingress LSRs, that is at the border of the network) and a partial waste of bandwidth dedicated to QoS traffic. A simulation was then run to study the performances of such mechanisms. Fig. 1 shows how to choose, according to the simulation's parameters, the proper signalling/bandwidth trade-off. In our simulation the use of caching and aggregation shows a dramatic reduction of the number of CR-LDP messages at about 2% of the original number.

6. REFERENCES

- [1] Rosen, E., Viswanathan, A. and R. Callon, "Multiprotocol Label Switching Architecture", RFC3031, January 2001.
- [2] Davie, B., Rekhter, Y., "MPLS Technology and Applications", Morgan Kaufmann, 2000.
- [3] Andersson, L. et al. "LDP Specification", RFC3036, January 2001.
- [4] Jamoussi, B. et al. "Constraint-based LSP Setup Using LDP", RFC3212, January 2002
- [5] Ash, J., Girish, M., Gray, E., Jamoussi, B., Wright, G., "Applicability Statement for CR-LDP", RFC3213, January 2002.
- [6] Swallow, G., "MPLS Advantages for Traffic Engineering", IEEE Communication Magazine, December 1999.
- [7] Vutukury, S., Garcia-Luna-Aceves, J.J., "A Simple MPLS-based Flow Aggregation Scheme for Providing Scalable Quality of Service", Proc. SPIE ITCOM 2001: International Symposium on The Convergence of IT and Communications, 19-24 August 2001, Denver, Colorado.