Computer Networks 127 (2017) 233–242

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Flow termination signaling in the centralized pre-congestion notification architecture

Frank Wetzels[a,*], Hans van den Berg[a,b,d], Joost Bosman[a], Rob van der Mei[a,c]

[a] Centrum Wiskunde en Informatica, Amsterdam, The Netherlands
[b] TNO, The Hague, The Netherlands
[c] Vrije Universiteit Amsterdam, The Netherlands
[d] Technische Universiteit Twente, Enschede, The Netherlands

ARTICLE INFO

Article history:
Received 14 February 2017
Revised 9 July 2017
Accepted 14 August 2017
Available online 30 August 2017

Keywords:
Centralized decision
Point pre-congestion notification
Signaling
Software defined networking
Flow termination

ABSTRACT

Pre-congestion notification (PCN) protects inelastic traffic by using feedback on network link loads on and acting upon this accordingly. These actions comprise to admission control and termination of flows. Two PCN architectures have been defined by IETF: the centralized and decentralized PCN architecture. The decentralized PCN architecture has received much attention in the literature whereas the centralized PCN architecture has not. In the decentralized architecture, feedback is sent from the egress nodes to ingress nodes, which then take and apply decisions regarding admission of new flows and/or termination of ongoing flows. Signaling occurs only between ingress and egress nodes.

In the centralized architecture these decisions are made at a central node, which requires proper signaling for action and information exchange between the central node and the egress and ingress nodes. This signaling has been suggested by other authors, but is not fully defined yet. Our contribution is twofold. We define signaling in the centralized PCN architecture focussing on flow termination, which completes the definition of the signaling in the centralized PCN architecture. Secondly, we run extensive simulations showing that the proposed signaling works well and that the performances of the centralized PCN and the decentralized PCN architectures are similar. Hence, it is expected that results from existing research on the effectiveness of decentralized PCN are also valid when the centralized PCN architecture is used.

1. Introduction

Currently, video and web traffic are major contributors to internet traffic. Web traffic is built upon an elastic transport protocol, mostly TCP which can adapt to congestion. Nowadays, also (non-real-time) video traffic like YouTube is increasingly delivered over TCP, which requires the video coding to be able to adapt in case of congestion. However, real-time video applications and VoIP use an inelastic protocol (e.g. UDP). Such protocol cannot adapt to congestion in the network and may suffer by packet loss, increased delay, greater jitter and reduced available bandwidth. This affects real-time applications like VoIP, VoD, IPTV and others. Which leads to a degradation of the quality of service (QoS) experienced by the users of real-time applications.

Pre-congestion notification (PCN) protects inelastic traffic by flow admission and flow termination [1] when certain criteria related to the network load are met [2,3]. Decisions to take actions are based on traffic measurements in the network and reporting upon these measurements. Traffic enters the PCN-domain at an ingress-node and leaves at an egress-node. While traffic flows through the network, passing internal nodes, traffic is classified against pre-defined PCN-related thresholds. Based on the amount of PCN marked traffic [4] a report is created at fixed time periods and sent to the decision making node. These reports may trigger admission control and flow termination decisions. When traffic leaves the network, marked traffic is administered for the next report to be sent. PCN can be applied in a centralized and decentralized architecture. In this paper we denote by cPCN and dPCN, the centralized and *de*centralized PCN architecture respectively. In dPCN, all egress-nodes send feedback to *ingress-nodes* which take and apply decisions on flows. In cPCN, all egress-nodes send feedback to a *central node*, the decision point (DP), which decides what to do upon such feedback.

After a decision is made, the ingress nodes need to get instructed what to do: Admit or block a new flow, i.e. admission control (AC), or terminate one or more existing flows, i.e. flow termination (FT). The signaling between the DP and the ingress-

* Corresponding author.
E-mail address: frank.wetzels@cwi.nl (F. Wetzels).

http://dx.doi.org/10.1016/j.comnet.2017.08.010

nodes has been suggested by other authors [5,6]. However, some essential components are missing in the signaling. This paper fills in the current gaps in cPCN signaling. In addition, extensive simulations have been carried out for both cPCN and dPCN as well as for a network without PCN in order to show the effectiveness of our proposed signaling. These simulations show that the proposed cPCN signaling works properly from a functional point of view, and that the performances of the cPCN and dPCN architectures are very similar. Hence, it is expected that results from existing research on the effectiveness of dPCN are also valid when cPCN is used. As in the aforementioned references [1–6], our specifications and simulations are based on 'traditional' networks assuming an interior gateway protocol and destination based forwarding. However, the cPCN signaling architecture fits very well to the centralized nature of the control architecture of emerging Software Defined Networks (SDN, see e.g. [7–9]) that (amongst others) takes care of flow routing in the data plane. Therefore, the outcome of our study also shows potential for enriching SDN with flow admission control and flow termination functionalities according to the cPCN approach. To the best of our knowledge such an extension of SDN has not yet been considered in the literature. The remainder of this paper is organized as follows. We start with a background on PCN and related work in Section 2. Section 3 highlights the proposed changes and additions to the signaling required in the cPCN. Section 4 describes these signaling modifications and additions in great detail for both admission control and flow termination. At the end of the section the identifiers and messages are defined in detail. In Section 5 the results of the simulations done in networks with cPCN, dPCN and without PCN are presented and discussed. Finally, discussions, conclusions as well as topics for future work are given in Section 6.

2. Background

The general architecture of PCN is given in [1]. If a new flow requests to enter the PCN-domain, it is decided whether or not this flow gets admitted to the PCN-domain (AC). This decision is based on the traffic load in the network. If an unusual event occurs in the network, for example a link failure, traffic gets rerouted and severe traffic overload on one or more links may happen. In such cases PCN may even decide to terminate one or more existing (previously admitted) flows (FT). The decision point (DP) decides whether a new flow gets admitted or blocked and what flows should be terminated, if applicable. In dPCN, each ingress node acts as DP for associated traffic, i.e. no central DP exists. In the cPCN, one node acts as DP. The DP does not take part in the data forwarding. The decision criteria for AC and FT are specified in [2,3] for the single marking (SM) and controlled load (CL) implementation respectively. In this paper we will focus on the signaling of the CL implementation in cPCN with one DP.

A brief overview of the research done on PCN is given below. In [10–14] the effectiveness of PCN is investigated in the context of a network with CBR traffic with on-off periods approximating different types of voice and video traffic. In particular, in [14] different PCN-based AC algorithms are considered and compared under various network load conditions. Reference [13] proposes a new measurement algorithm (sliding window) for AC based on bandwidth metering. In [15] an autonomous AC algorithm is proposed optimized for bursty traffic, which adapts itself based on previous measurements. Performance and parameter sensitivity analysis is done in [16] for both the SM and CL in dPCN. In [17] an summary is given of many aspects of PCN including the working, benefits, signaling and limitations of PCN in general.

We will now focus on the signaling in cPCN, in particular the associated signaling aspects. To determine whether AC and/or FT is required, the DP needs feedback from the egress nodes. The feedback is generated per aggregate at fixed time intervals by egress nodes and sent to the DP. An ingress-egress-aggregate, aggregate in short, is a set of flows which travel in the network from an ingress node to an egress node. The DP needs to exchange information with ingress nodes as to what the actual aggregate rate is, inform on whether to admit or block flows and to inform the ingress node(s) which ongoing flow(s) need to be terminated, if the FT criterium is met. The egress nodes need to send feedback to the DP which should contain information on the load per aggregate.

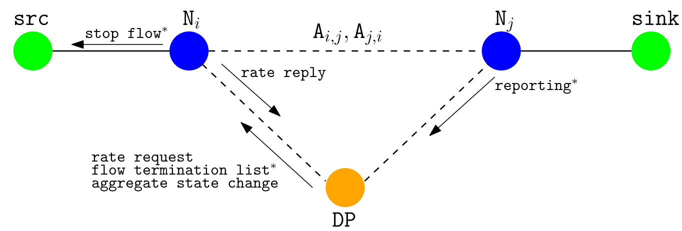

Fig. 1. Signaling data flow in cPCN. An '*' indicates a change to the current definition or a new definition.

On the signaling in a PCN-domain, P. Eardley [1] refers to related work that consider specific signaling protocols or frameworks like next steps in signaling (NSIS, [18]), resource reservation protocol (RSVP, [19]) and extensions to RSVP [20]. In [2], signaling is considered out of scope and refers to [20] as well. NSIS mainly focuses on protocols for signaling that follow the same paths along which the user-data flows, i.e. path-coupled signaling. NSIS considers the path-*de*coupled signaling briefly. In SDN and cPCN, all signaling is decoupled from the data path since all signaling happens between SDN switches and the SDN controller. In [5], requirements for signaling in a PCN-domain are described. Karagiannis et al. [5] restricts to feedback signaling between egress-nodes and DP and the signaling between DP and ingress-node on the aggregate-rate request. The signaling between DP and ingress-nodes on which flows to terminate and how to stop a source from sending a current (to be terminated) flow is not specified. For that, a reference is made to the common open policy service architecture (COPS, [21]) and the diameter based protocol (DBP, [22]) as a basis for a full signaling architecture. In [6] a signaling protocol, regular-check-based flow termination (RCFT), is proposed using RSVP as a carrier. It fills in the gap in the FT-communication between egress and ingress nodes. However, RCFT is focused on dPCN. In [17] the path-decoupled signaling in cPCN is discussed. However, it does not define the actual signaling in case of termination of flows. In this paper, we will propose signaling in case of flow termination and make an addition to the reporting. Simulation is used to check the functional correctness of these extensions and evaluate their performance.

3. Signaling in the cPCN

In this section the signaling between ingress-egress nodes, ie-nodes in short, and DP is considered, i.e. PCN signaling in the cPCN. The following components will be introduced: the *flow-rate*, the *flow-termination list* and the *flow-off* signal.

Refer to Fig. 1. The focus will be on two ie-nodes and one DP. This small network with one DP is no restriction as for every edge-node in the network the signaling below still applies. Considering multiple DPs would introduce other issues, like synchronization between DPs and the placement of DPs as well. These issues would distract our focus from the signaling. Between the ie-nodes two unidirectional aggregates exist. By $A_{i,j}$, we refer to the unidirectional aggregate from ie-node N_i to ie-node N_j The DP will not be part of any data-path, i.e. no aggregate will flow through the DP.