Computer Networks 95 (2016) 35–50

Contents lists available at ScienceDirect

Computer Networks

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

Routing optimization for IP networks with loop-free alternates

Matthias Hartmann[a], David Hock[a], Michael Menth[b,*]

[a] Institute of Computer Science, University of Würzburg, Würzburg, Germany
[b] Department of Computer Science, University of Tübingen, Tübingen, Germany

ARTICLE INFO

Article history:
Received 20 June 2014
Revised 1 October 2015
Accepted 4 November 2015
Available online 21 December 2015

Keywords:
IP fast reroute
Loop-free alternates
Routing optimization
Coverage
Maximum link utilization
Traffic engineering

ABSTRACT

Loop-free alternates (LFAs) have been developed for fast reroute (FRR) in intradomain IP networks. They are simple, standardized, and already offered by several vendors. However, LFAs have two major drawbacks. They often cannot provide failure protection against all single link or node failures in spite of physical connectedness, and some LFAs cause routing loops in scenarios with node or multiple failures.

LFAs may be applied for various reasons that we call applications in this work. We propose several definitions for LFA coverage that quantify the application-specific utility of LFAs available in the network. The availability of LFAs and whether they can cause routing loops heavily depend on the IP routing which is determined by the choice of administrative IP link costs. To maximize the benefit of LFA usage, we optimize the IP link costs using LFA coverage as objective function. We demonstrate the feasibility and effectiveness of that approach in several test networks, and show that the choice of the right optimization function is crucial to maximize LFA coverage. However, maximizing LFA coverage can lead to significant traffic imbalance and may result in high link loads. Therefore, we suggest Pareto-optimization and demonstrate that resulting link costs can lead to both high LFA coverage and low link loads.

1. Introduction

In IP networks, failures occur on a regular basis and often last only for a short time [1]. The distributed IP rerouting process is simple and robust [2], but it may be too slow for applications and services that require continuous network availability [3]. Recently, fast reroute (FRR) mechanisms have been proposed for IP networks [4]. With IP-FRR, a router can detour traffic around a failure location immediately after it has detected that the regular next-hop is no longer reachable. This reduces the time during which packets are lost from several seconds down to less than 50 ms. Then, regular IP rerouting is triggered. Therefore, the traffic affected by the failure is forwarded by IP-FRR mechanisms only until the rerouting process completes or the failure disappears.

The only IP-FRR mechanism that is already standardized by the IETF and implemented in new routers, e.g., current versions of Cisco IOS and Juniper OS, is the loop-free alternates (LFAs) concept [5]. An LFA is an alternate next-hop to which certain traffic can be sent without creating any loops so that this traffic reaches its destination over an alternative path. When the regular next-hop for a certain destination is no longer reachable by a router, it can deflect traffic to this destination over the LFA. LFAs do not require any signaling, they do not require changes to the basic IP routing protocol, and they do not require tunneling. These features facilitate incremental as well as partial

* Corresponding author. Tel.: +49 7071 2970505.
E-mail address: menth@uni-tuebingen.de (M. Menth).

http://dx.doi.org/10.1016/j.comnet.2015.11.009
1389-1286/

deployment, even in a multi-vendor network, and make LFAs a very attractive solution. However, LFAs have also two disadvantages. First, nodes may not have LFAs for all destinations [6–8] so that some traffic cannot be protected against single link or node failures although the network topology has alternate working paths. Second, some LFAs may cause extra-loops in case of node or multiple failures. An extra-loop is a forwarding loop caused by LFAs where packets loop between two or more nodes. This can even overload links and routers that are otherwise unaffected by the failure.

There are various incentives for the use of LFAs in IP networks. We call them applications and consider several of them. We argue that the utility of available LFAs depends on the application and measure the utility by application-specific LFA coverages. Some examples:

- LFA coverage can be measured by the fraction of destinations that each node can protect by LFAs, averaged over all nodes. This is an intuitive definition that nicely reflects the availability of LFAs in a network and was used for that purpose in most existing studies on LFAs. However, it does not relate to any specific application.
- One goal of IP-FRR is to reduce traffic loss between failure detection and the completion of the rerouting process. This is reflected by the fraction of the traffic that is lost due to missing LFAs, averaged over all considered failures. We use that as indirect measure for LFA coverage.
- Network providers can sell improved availability guarantees if traffic is protected by LFAs on its entire path so that only marginal traffic is lost in case of a failure. Thus, the LFA coverage may be quantified by the fraction of traffic for which the entire path can be protected by LFAs.
- If all flows carried over a link can be protected by LFAs, that link may fail without losing any traffic after LFA activation. As a consequence, IP rerouting may be delayed when such a link fails and graceful reconvergence techniques [9–12] can be utilized that prevent micro-loops. For short-lived link failures or maintenance operations, IP rerouting that can lead to routing instabilities and micro-loops, may be avoided even twice: once when the link goes down and once when it comes up again. For these applications, the LFA coverage may be expressed by the fraction of links for which all traffic carried under failure-free operation can be protected by LFAs.

We further diversify the definitions of LFA coverage with regard to the types of LFAs that may be used: all LFAs or only those that cannot create extra-loops. The relevance of avoiding temporary extra-loops is certainly application-specific.

The availability of LFAs and the LFA coverage obviously depend on the network topology and the routing. Thus, LFA coverage may be increased by changing the topology: additional (physical or virtual) links may be installed which provide LFAs that can be used during failures [13,14]. LFA coverage can also be increased by changing the routing by configuration of appropriate administrative link costs that determine the path layout in IP networks [15,16].

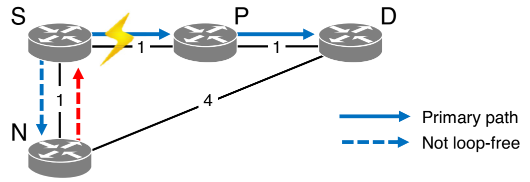

Fig. 1. Neighbor N cannot be used as LFA because it does not meet the loop-free condition.

In this work, we investigate the different definitions of LFA coverage in test networks with uniform link costs. We further apply these definitions as objective functions to optimize link costs in order to maximize LFA coverage. We show that this approach is feasible by achieving significant improvements in LFA coverage. However, tweaking link costs influences not only LFA coverage but also traffic distribution within the network. We show that maximizing LFA coverage can lead to significantly increased link loads both under failure-free conditions and after rerouting in failure cases so that traffic may be lost due to overload. That is not acceptable since these phases persist longer than the short rerouting interval for which LFAs reduce packet loss. Hence, maximization of LFA coverage can be counterproductive. To fix that problem, we propose Pareto-optimization to generate a set of link costs that are Pareto-optimal with regard to LFA coverage and maximum link loads. Some of these link costs lead to relatively high LFA coverage and relatively low maximum link loads so that a network administrator can choose appropriate ones to configure the network.

The remainder of this paper is structured as follows. Section 2 explains LFAs and Section 3 gives an overview of related work. Section 4 discusses various applications of LFAs that require different definitions of LFA coverage, and the potential of routing optimization is illustrated. Section 5 shows that there is a tradeoff between high LFA coverage and low link loads and suggests Pareto-optimization to find good compromises. Finally, Section 6 summarizes this work. A table with acronyms and notation is provided in Table 10 of the Appendix.

2. Loop-free alternates

LFAs provide fast protection for IP networks using link state routing protocols. They are intended to be used by a node immediately after it has detected a failure until the failure disappears or until IP rerouting has converged. In this section we review the definition of LFAs [5]. As general LFAs may cause extra-loops under some conditions, we define three sets of LFAs that avoid extra-loops to a different extent.

2.1. General or link-protecting LFAs

We consider a *source node* S and a *next-hop* P on a shortest path toward *destination* D, just like in Fig. 1, but with a link cost less than 3 for the link from N to D. In this scenario, another *neighbor node* N of S can be used by S as LFA to D for the potential failure of the link $S \rightarrow P$