Computers & Operations Research 38 (2011) 1805–1815

Contents lists available at ScienceDirect

Computers & Operations Research

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

Linear programming models for traffic engineering in 100% survivable networks under combined IS–IS/OSPF and MPLS-TE

D. Cherubini [a,*,1], A. Fanni [b], A. Mereu [b], A. Frangioni [c], C. Murgia [d], M.G. Scutellà [c], P. Zuddas [e]

[a] *Tiscali International Network, loc. "Sa Illetta", Cagliari, Italy*
[b] *DIEE, Università di Cagliari, P.za d'Armi 1, Cagliari, Italy*
[c] *Dipartimento di Informatica, Università di Pisa, Largo B. Pontecorvo 3, Italy*
[d] *Tiscali Italia S.p.A., loc. "Sa Illetta", Cagliari, Italy*
[e] *DIT, Università di Cagliari, P.za d'Armi 1, Cagliari, Italy*

ARTICLE INFO

Available online 25 February 2011

Keywords:
Routing problems
LP models
Robustness

ABSTRACT

This paper concerns the problem of minimizing the maximum link utilization of IP telecommunication networks under the joint use of traditional IGP routing protocols, such as IS–IS and OSPF, and the more sophisticated MPLS-TE technology. It is shown that the problem of choosing the optimal routing, both under working conditions and under single link failure scenarios, can be cast as a linear program of reasonable size. The proposed model is validated by a computational experimentation performed on synthetic and real networks: the obtained results show that the new approach considerably reduces the maximum link utilization of the network with respect to simply optimizing the IGP weights, at the cost of adding a limited number of label switched paths (LSPs). Optimizing the set of IGP weights within the overall approach further improves performances. The computational time needed to solve the models matches well with real-time requirements, and makes it possible to consider network design problems.

1. Introduction

Communication and computer networks have high deployment and maintenance costs, which makes development of optimization models for their design and routing an important subject to both researchers and practitioners [1,2]. This is even more true as the continuous development of new technologies keeps offering new avenues for optimization, while requiring to re-think established operations. Several different forms of network design problems exist depending on the specific details of the network under consideration: the *layer* (physical, IP,…), the corresponding *routing protocols*, the routing and deployment costs, and the *survivability requirements* specifying under which set of conditions the network should be guaranteed to remain operational. While very general forms of design and routing problems can be devised [3], being able to solve them within the time limits dictated by the application environments crucially depends on exploiting the specific technological characteristics of the network at hand. Many of these problems are hard to solve, both in theory and in practice. For instance, optimal routing under the available IGP (interior gateway protocol), such as intermediate system to intermediate system (IS–IS) or open shortest path first (OSPF), requires finding a set of weight links (integers in the range $[1, 2^{24}]$) that completely determine the routing of the demands; even with no survivability assumptions this is a $\mathcal{NP}$-hard problem that cannot be approximated with a factor better than 3/2 [4]. Thus, even simple routing on current IP networks is a difficult problem; fortunately, new traffic engineering (TE) techniques have been developed to enable Internet Service Providers (ISPs) to route the traffic along the network. In particular, multiprotocol label switching (MPLS-TE) networks enable ISPs to move traffic from congested links to less loaded areas of the network setting up constrained label switched paths (LSPs). However, the use of a "full mesh" MPLS-TE, where each origin–destination pair gets its dedicated LSP (possibly more than one to take into account different kinds of traffic) is not scalable, since it typically requires $O(n^2)$ paths which is not considered to be practical. Hence, one should rather search for an optimal set of complementary LSP tunnels to be used in combination with the IGP protocol.

This is especially important as, besides optimizing network resources utilization in the *nominal* case, these techniques can be used to implement *network survivability techniques* which allow resilience to equipment failure or congestion. Several different

* Corresponding author. Tel.: +353 1 8864602.
E-mail addresses: davide.cherubini@alcatel-lucent.com (D. Cherubini), fanni@diee.unica.it (A. Fanni), anna.mereu@diee.unica.it (A. Mereu), frangio@di.unipi.it (A. Frangioni), cmurgia@tiscali.com (C. Murgia), scut@di.unipi.it (M.G. Scutellà), zuddas@unica.it (P. Zuddas).
[1] The author is now with Alcatel-Lucent Bell Labs Ireland, Blanchardstown Ind. Park, Blanchardstown, Ireland.

0305-0548/$ - see front matter © 2011 Elsevier Ltd All rights reserved.
doi:10.1016/j.cor.2011.02.019

optimization models dealing with different forms of survivability have been developed [5–18]; in this paper we focus on *single link* failures, which represent by far the most common form of unplanned failure in real networks [19]. To survive a (single link) failure, *restoration routing* is performed; this can be *path* (end-to-end) or *link* (local) restoration [20]. In the former scheme, if a failure is present along a path, the source node detects it and activates backup paths; due to its global nature, this may cause very large restoration latencies. In the link restoration scheme, a backup path is found to protect the failed link, rather than the entire path; although this may use more resources, it is attractive in that it can meet strict restoration latency requirements.

Furthermore, the local restoration scheme fits very well with a specific combination of network technology as far as the resulting optimization models are concerned. In particular, we will show that if restoration paths are all chosen by the IS–IS/OSPF protocol, *for a fixed set of weight links* the survivable routing problem using a combination of the available routing technologies (IS–IS/OSPF and MPLS-TE) can be modeled as a relatively small-size linear program (LP). While the assumption of fixed weight links may appear to be a rather strong one, our computational results show that the model provides very good resources utilization with any reasonably chosen set of IS–IS weights, while dictating the use of a small set of LSP tunnels and requiring a comparatively short running time to be solved by off-the-shelf methods. This is somewhat surprising, as several routing problems under realistic conditions are difficult, especially if survivability has to be taken into account, as discussed in Section 3. Indeed, while several optimization models under the MPLS-TE technology have been presented in the literature [20–26], the majority of these papers either only consider the normal (or working) condition, or deal with survivability constraints as a two phase design process, where the working paths are firstly chosen to optimize the routing in working conditions, and only then restoration (or backup) paths are selected that protect the links along the primary LSP. To the best of our knowledge, no paper considers at the same time optimal routing and survivability in the specific (but very reasonable in practice) set of technological requirements described here, which fortunately happen to make it easy to solve. This is interesting in that the efficient solution of the routing problem is usually a crucial step in the far more complex network design problems in telecommunication networks [27,9,2,1].

The outline of the paper is as follows. Section 2 first presents some relevant backgrounds on various aspects of routing problems, then summarizes related works dealing with IGP and MPLS-TE optimization approaches. In Section 4 we describe the routing optimization models, while in Section 5 we demonstrate the suitability of the proposed optimization models by tests and comparisons on synthetic and real networks. Finally, Section 6 presents our conclusions and directions for future research.

2. Background

2.1. Routing protocols

IP backbone networks employ IGP routing protocols, such as IS–IS and OSPF, in order to compute the routes from each origin node to each destination node of the traffic demands. This is done by assigning a suitable metric value to each link of the network, and then routing along the shortest paths w.r.t. the selected metrics (weights). When equivalent cost multiple paths (ECMPs) are enabled, the traffic flow is equally split among the alternative shortest paths linking each origin–destination pair. Generally, the traffic splitting mechanism is *per packet round robin*, where each packet matching a given destination is forwarded toward the corresponding egress node using the least recently used equal cost path.

The introduction of the multi-protocol label switching traffic engineering (MPLS-TE) technology has improved the dataflow management due to the traditional routing protocols. In fact, unlike IGP routing protocols, MPLS-TE networks can support destination-based and explicit routing simultaneously. Furthermore, MPLS-TE provides better handling of congestion and failures by supplying mechanisms to quickly find an alternate path if the primary path is no longer available. This fast re-route (FRR) capability is critical for allowing service providers to offer high availability and high revenue service level agreements. However, FRR requires that the operator explicitly sets secondary paths for all the possible failure scenarios, which can be excessively time consuming; besides, choosing appropriate reserve paths is non-trivial. An interesting alternative is therefore to rely on the link restoration approach, which is technically obtained by using "Implicit LSP" as reserve paths; this causes the commodities affected by the fault to be re-routed using the IGP (IS–IS/OSPF), thereby negating the need to explicitly choosing the reserve paths.

3. Related works

Several strands of research are related to the results in this paper.

The first concerns optimal routing using IS–IS/OSPF protocols. This can be cast in different ways [28], and several (mostly heuristic) approaches have been proposed to solve it such as Tabu Search (TS) [29], genetic algorithms [30,31] and tailored heuristics [32,33]. An extended comparison of different heuristic approaches, comprising local search, simulated annealing (SA) and Lagrangian relaxation, is presented in [34], both for the case where ECMPs are allowed and for that where only a unique path for demand is allowed. The results of these papers show that optimal IS–IS/OSPF metrics not only are cumbersome to find, but also do not perform as well as an optimal MPLS-TE routing.

Other papers have addressed the problem of determining an optimal MPLS routing according to different objective functions. In particular, in [22] a TS algorithm is proposed to find a layout of MPLS-TE paths with the minimum number of hops, whereas in [23] the MPLS routing problem is modeled in terms of an off-line multiobjective mixed integer linear programming (MILP) model that looks for the best trade-off between the minimal routing delay, the optimal load balancing, and the minimal splitting of traffic trunks. In [35,36] a multipath adaptive traffic engineering mechanism is introduced which aims at avoiding the network congestion by adaptively balancing the load among multiple paths based on measurements and analysis of the path level. In [37,38] an online traffic engineering approach is proposed which aims to determine minimum interference LSP tunnels, defined as paths that maximize the weighted sum of flows between all the other node pairs (called WSUM-MAX problem): a heuristic online procedure is proposed which allows a good LSP acceptance rate and re-routing performance, thus limiting the well-known online traffic engineering drawbacks related to unavailability of bandwidth at the request time.

Relatively few papers deal with combined IGP/MPLS routing, and most of these employ a two phase (hierarchical) approach by optimizing separately the two technologies. In [39] the routing problem is formulated as a service differentiated model where the IP flows are classified (based on experience) into low bandwidth demanding (LBD) and high bandwidth demanding (HBD) flows at the ingress of the network; LBD flows are routed along the IGP shortest paths, while HBD flows over bandwidth-guaranteed LSP