

Differential delay constrained multipath routing for SDN and optical networks

Rodolfo Alvizu^{a,b,1} Guido Maier^{a,2} Massimo Tornatore^{a,3}
Michał Pióro^{c,d,4}

^a *Dipartimento di Elettronica, Informazione e Bioingegneria (DEIB),
Politecnico di Milano, Italy*

^b *Department of Electronics and Circuits, Simon Bolivar University, Venezuela*

^c *Institute of Telecommunications, Warsaw University of Technology, Poland*

^d *Department of Electrical and Information Technology, Lund University, Sweden*

Abstract

In multipath routing, maximization of the cardinality K of the disjoint-path set for a given source and destination assuming an upper bound on the differential delay D is one of the key factors enabling its practical applications. In the paper we study such an optimization problem for multipath routing involving maximization of K under the D constraint as the primary objective, and then minimization of the average end-to-end transfer delay for the fixed (maximum) K under the same D constraint. The optimization approach is iterative, based on solving an inner mixed-integer programming subproblem to minimize the delay for a given value of K and D . In order to increase the solution space, we consider the strategy of allowing controlled routing loops. Such a technique is implementable in software defined networks and optical networks. We present numerical results illustrating the gain achieved by using controlled loops in comparison with the traditional loop-free approach.

Keywords: multipath routing, differential delay, loops, integer programming

¹ Email: rodolfoenrique.alvizu@polimi.it

² Email: guido.maier@polimi.it

³ Email: massimo.tornatore@polimi.it

⁴ Email: mpp@tele.pw.edu.pl

1 Introduction

Multipath (MP) routing is a network functionality controlling splitting of the data flow from a source to a destination among multiple paths, and reconstructing it at the destination as a single flow, prior to being delivered to an upper layer. In SDH networks, for instance, the MP VCAT technique allows for a better utilization of network resources [3]. In optical transport networks (OTN), MP can be exploited to drastically decrease the amount of bandwidth reserved for protection [7]. In another context, MP is used by the Multipath Transport Control Protocol (MPTCP) [5] to augment the throughput of the TCP-based applications in a transparent way, i.e., without modifying the applications and yet preserving backward compatibility with TCP.

Today's networks generally offer MP capability, but the common networking practice seldom exploits this. One of the major obstacles is the data reconstruction operation at the destination: if significant differences in the delays occur between the paths of the MP set, the reconstruction buffer has to be increased in size and the reordering task becomes time consuming. Consequently, the quality of the service delivered over the MP connection starts degrading, causing users experiencing unacceptable waiting times to receive the data (e.g., in conversational or gaming applications) or – even worse – an unexpected bursty-mode operation (e.g., in video streaming). Another issue is related to routing: MP improvements in terms of throughput, load balancing, reliability and protection bandwidth are all fully achievable only provided that the paths of the MP are link disjoint. For example, it was proven that the TCP performance can be enhanced by MPTCP only if physical path-disjointness is enforced [4]. In a highly loaded network the throughput of the MP connection will increase as the number of paths K grows.

The problem of Differential Delay Routing (DDR) was first studied by [1] and [3] for the Ethernet over SONET architectures. The DDR problem was defined as follows: given a graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, find K paths of unit capacity from source S to destination T such that their differential delay D is upper bounded by a given constant Δ . D is defined as the maximum end-to-end delay difference between the K paths of the MP set. Link disjointness was not considered in that DDR problem. More recently, the work presented by Sheng *et al.* [7] included link disjointness as a constraint. However, the authors relied on heuristics to solve the problem, as DDR (and consequently DDR with disjointness) is NP-hard [3].

In this paper we extend the previous work by redefining the problem as follows: (a) find a set of K link-disjoint paths from source S to destination

T that maximizes K such that D is upper-bounded by a given value Δ ; (b) find the set of K^0 link-disjoint S - T paths minimizing the average end-to-end delay L (where K^0 is the maximum found in phase (a)) under the same D constraint. We call this the 3D problem (Disjoint-Differential-Delay).

In fact, when looking for maximum K in phase (a) we could assume some fixed upper bound on L as well. Otherwise, looking for a large cardinality set of multiple paths with similar delays could easily result in selecting very long paths. This would have a negative impact on the quality of service, as it would slow down the applications (especially when automatic congestion control is adopted, as with TCP or MPTCP), and at the same time would increase network congestion.

To reduce the complexity of the 3D problem, we will adopt an iterative procedure involving a subproblem solved by means of a mixed-integer programming (MIP) formulation, as described in Section 2.

In general, enforcement of the D constraint in the 3D problem can result in appearance of routing loops associated with the paths. Such cycles are usually undesired in networking but in the MP context they can compensate for delay differences without adding buffering capabilities at the destination nodes or transit nodes of a path, as proposed in [2].

Some network technologies, such as software defined networks [6], MPLS and optical networks, are potentially able to handle routing loops without generating “broadcast storms” (infinite loops) or routing failures. Therefore, a 3D formulation that allows for loops is also presented.

The comparison of the numerical results obtained by the loop-free model (3D-LF) with the results of the model which allows loops (3D-LIP) exhibits an advantage of 3D-LIP. This justifies the interest for the 3D-LIP case. In the paper we focus on the optimization aspects and do not discuss practical implementation of 3D-LIP in real networks.

2 Optimization models

Consider a network represented by an undirected graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, where \mathcal{V} and \mathcal{E} are the sets of nodes and (undirected) links, respectively. Each link $e \in \mathcal{E}$ is associated with two oppositely directed arcs e' and e'' joining its nodes. The set of all such arcs is denoted by \mathcal{A} (note that the resulting directed graph $\mathcal{G}(\mathcal{V}, \mathcal{A})$ is bi-directed). The source and destination nodes of arc a will be denoted by $s(a)$ and $t(a)$, respectively. The weight w_a of arc $a \in \mathcal{A}$ represents its delay. The sets of arcs $\delta^+(v)$ and $\delta^-(v)$ represent the outgoing star of arcs from node $v \in \mathcal{V}$ and the incoming star of arcs directed to v , respectively.

Problem 3D consists in finding an MP connection between a given source S and destination T node pair in graph \mathcal{G} that complies with given differential delay Δ and average end-to-end delay Λ upper bounds, and as first objective maximizes the number of disjoint paths K and then minimizes the average end-to-end delay L .

The MP connection is a set of K link-disjoint paths \mathcal{P} from S to T in the directed graph $\mathcal{G}(\mathcal{V}, \mathcal{A})$. The paths in \mathcal{P} are denoted by $P_k, k \in \mathcal{K}$, where $\mathcal{K} = \{1, 2, \dots, K\}$. The delay W_k experienced by the k^{th} path in \mathcal{P} is given by $W_k = \sum_{a \in P_k} w_a, k \in \mathcal{K}$. The average end-to-end delay of the MP connection \mathcal{P} is given by $L = \frac{1}{K} \sum_{k \in \mathcal{K}} W_k$.

In order to simplify the formulation and to make 3D problem tractable, we introduce an iterative procedure for maximizing K for a given value of Δ :

- 1: Set differential delay Δ and average end-to-end delay Λ upper bounds.
- 2: Initialize the number of paths $K = 2$.
- 3: Solve the optimization problem 3D* with K as an input parameter.
- 4: If no solution is found or the optimal L is greater than Λ then go to 5; else $K = K + 1$ and go to 3.
- 5: If $K = 2$ then reject the MP connection; else return the result obtained for $K = K - 1$ and stop.

Consecutive instances of problem 3D* (see below) are solved for consecutive increasing K (starting with $K = 2$) until the problem becomes infeasible.

The formulation of the optimization (sub)problem 3D* is as follows. An MP connection request is described by the quadruple (S, T, Δ, K) , where S, T are the source and destination nodes, Δ is the differential delay upper bound, and K is the assumed number of disjoint paths.

A node-link formulation of the 3D* problem is given in (1). The formulation uses binary arc-flow variables x_{ak} ($x_{ak} = 1$ if arc $a \in \mathcal{A}$ is used by path $P_k, k \in \mathcal{K}$; 0, otherwise) and the nonnegative variable h_{vk} (h_{vk} is the hop count from the source to a node $v \in \mathcal{V}$ for each path $P_k, k \in \mathcal{K}$).

In the formulation, objective (1a) minimizes the cumulative delay of the MP connection. The flow conservation constraint (1b) makes the flows x_{ak} represent K paths from source S to destination T . Constraint (1c) forces link disjointness of the paths. The differential delay constraint is formulated in (1d)-(1e). The last constraint (including a “big M” parameter) (1f) ensures that loops, otherwise advantageous in optimal solutions from the viewpoint of (1d)-(1e), cannot appear. In consequence, any feasible solution of (1) will be loop-free. Formulation (1) will be called 3D*-loop-free (3D*-LF in short).

$$\mathbf{3D^*-LF:} \text{ minimize } \sum_{a \in \mathcal{A}} \sum_{k \in \mathcal{K}} w_a x_{ak} \quad (1a)$$

$$\sum_{a \in \delta^+(v)} x_{ak} - \sum_{a \in \delta^-(v)} x_{ak} = \begin{cases} 0, & v \in \mathcal{V} \setminus \{S, T\} \\ 1, & v = S \\ -1, & v = T \end{cases}, \quad v \in \mathcal{V}, k \in \mathcal{K} \quad (1b)$$

$$\sum_{k \in \mathcal{K}} x_{e'k} + x_{e''k} \leq 1, \quad e \in \mathcal{E} \quad (1c)$$

$$\sum_{a \in \mathcal{A}} w_a x_{ak} - \sum_{a \in \mathcal{A}} w_a x_{am} \leq \Delta, \quad k, m \in \mathcal{K}, k < m \quad (1d)$$

$$\sum_{a \in \mathcal{A}} w_a x_{am} - \sum_{a \in \mathcal{A}} w_a x_{ak} \leq \Delta, \quad k, m \in \mathcal{K}, k < m \quad (1e)$$

$$h_{t(a)k} - h_{s(a)k} \geq 1 - M(1 - x_{ak}), \quad a \in \mathcal{A}, k \in \mathcal{K} \quad (1f)$$

$$x_{ak} \text{ binary, } h_{vk} \geq 0. \quad (1g)$$

Note that K is an input parameter of the formulation, and not an objective.

Now we introduce an important and novel variation of formulation (1). As we have already mentioned, formulation (1) eliminates loops due to (1f), both in-path loops (like F-G-H in path S-F-G-H-F-T in Figure 1a) and isolated loops (disjoint with the main part of an $S - T$ path). Although the latter paths are clearly not allowed, the former could be effectively used to compensate path delays as required by the D requirement. Thus, now we assume that the in-path routing cycles are allowed and manageable by the network in order to accomplish differential delay compensation. The appropriate formulation, called 3D*-LIP (loops-in-paths), is similar to the formulation of 3D*-LF (1), the differences are exposed in the following.

As 3D*-LF, the 3D*-LIP formulation uses binary flow variables x_{ak} to specify the path-set. Now, however, we use two extra sets of auxiliary non-negative continuous variables. The first set, z_{avk} ($z_{avk} = 1$ if arc $a \in \mathcal{A}$ is used by an artificial flow of value 1 from node $v \in \mathcal{V}$ on path $P_k, k \in \mathcal{K}$, to T ; 0, otherwise). The second set, r_{vk} ($r_{vk} = 1$ if node $v \in \mathcal{V}$ belongs to path $P_k, k \in \mathcal{K}$; 0, otherwise) is used to define the values of artificial flows.

The 3D*-LIP formulation shares the objective function (1a) and the first four constraints (1b)-(1e) with 3D*-LF (1). However, following [8], the loop avoidance constraint of 3D*-LF (1f) is exchanged by inequalities (2a)-(2b) and equalities (2c) in order to allow in-path loops but avoid isolated cycles. Artificial flow z_{avk} is initiated at each node $v \in \mathcal{V} \setminus \{S, T\}$ if $x_{ak} = 1$ and $a \in \delta^+(v)$. Constraints (2a), (2b) and (2c) define an artificial flow connecting node v to the destination node T . Clearly, equations (2c) represent the flow conservation constraints in a graph induced by the links in path P_k , i.e., in the graph $(\mathcal{V}, \{a \in \mathcal{A} : x_{ak} = 1\})$.

$$z_{avk} \leq x_{ak}, \quad a \in \mathcal{A}, v \in \mathcal{V} \setminus \{S, T\}, k \in \mathcal{K} \quad (2a)$$

$$r_{vk} \geq x_{ak}, \quad a \in \delta^+(v), v \in \mathcal{V} \setminus \{S, T\}, k \in \mathcal{K} \quad (2b)$$

$$\sum_{a \in \delta^+(u)} z_{avk} - \sum_{a \in \delta^-(u)} z_{avk} = \begin{cases} 0, & u \in \mathcal{V} \setminus \{v, T\} \\ r_{vk}, & u = v \end{cases}, \quad u \in \mathcal{V} \setminus \{S\}, v \in \mathcal{V} \setminus \{S, T\}, k \in \mathcal{K} \quad (2c)$$

$$x_{ak} \text{ binary}, \quad z_{avk}, r_{vk} \geq 0. \quad (2d)$$

3 Numerical study

In this section we compare optimal solutions of 3D-LF and 3D-LIP in terms of the MP cardinality K and its average end-to-end delay L , considering Δ as a parameter. For simplifying the presentation, we do not impose any bound on L used in Step 4 of the procedure in Section 2.

The sample weighted graph shown in Figure 1a illustrates how LIP can improve the 3D problem. For $\Delta = 0$, 3D-LF is infeasible, while 3D-LIP successfully finds two paths with equal delay by exploiting the LIP path S-F-G-H-F-T for delay compensation. For $\Delta = 1$ the optimum for both formulations is $K = 2$. However, the L of the solution provided by 3D-LIP is smaller.

Now we compare the two formulations on the 14-node network depicted in Figure 1b. The delay of each link w_e is random with uniform distribution in $\{1, 2, \dots, 5\}$. Figure 2 shows the results for $\Delta = 0, 1, \dots, 6$, averaged over 100 instances of such random delay settings. (We have skipped confidence intervals as they do not influence the general picture.)

As expected, for both formulations the cardinality of MP connection is strictly increasing with Δ . For the source S and destination T node pair in the network of Figure 1b the maximum number of link-disjoint paths is equal to 6, imposing the upper bound for the maximization of K . Although the gain obtained by allowing LIP is not so significant, Figure 2a shows how 3D-LIP improves the solution subset of Δ values. In particular, for $\Delta = [4, 5, 6]$, 3D-LIP always reaches the upper bound of $K = 6$ (while 3D-LF does not), with a small impact on L as compared with 3D-LF (see Figure 2b).

It can be noticed in Figure 2b that L is almost always decreasing with Δ and that L is forced to considerably increase as the requirement on D gets more stringent. In particular, maximum L is observed for $\Delta = 0$. This reveals




Fig. 1. (a) Comparison of 3D-LIP and 3D-LF. (b) A test network

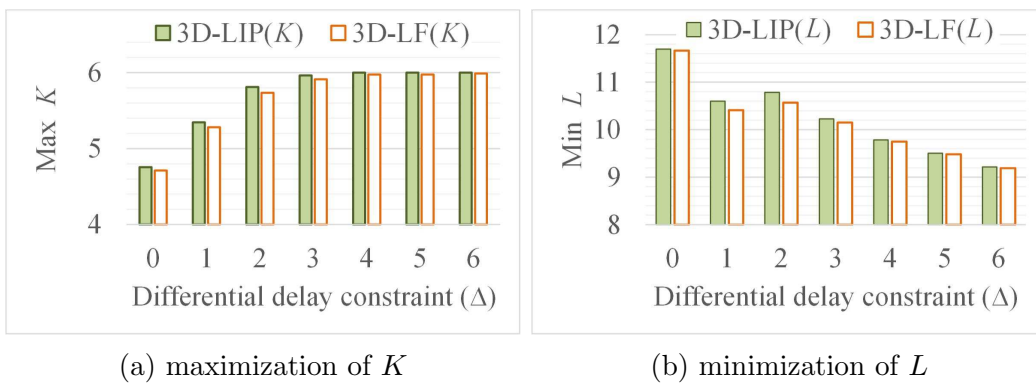


Fig. 2. Comparison of 3D-LIP and 3D-LF in a 14-node test network.

a clear trade-off between delay minimization and delay equalization in MP communications. 3D-LIP(L) is slightly larger than 3D-LF(L) due to those cases when 3D-LIP finds more paths than 3D-LF.

4 Concluding remarks

In the paper we have proposed an iterative optimization approach that simplifies the 3D (Disjoint-Differential-Delay) formulation and makes the problem tractable. We have shown how controlled in-path routing cycles broaden the solution space of the problem and can be used to equalize path delays as required by the D constraint. The optimization procedure 3D-LIP that allows for such loops always finds better (or equally good) solutions than the traditional loop-free approach (3D-LF) when maximizing the number K of disjoint paths in order to increase the throughput of the MP connection. For both formulations there is a trade-off between the differential delay and the average end-to-end delay of the MP connection. This trade-off must be resolved, depending on the application, by proper setting of D and L upper bounds.

This work represents an initial step in considering important issues of the MP networking, as we have focused only on the delay measures (D and L) and

assumed unlimited network capacity. In further steps, we will include such additional features as finite link capacity, delay variations over time, node disjointness, and concurrent routing of several MP demands. This additional features will make the problem even harder to solve.

The presented model can be applied to software defined, MPLS, and optical networks. For instance, an optimal routing solver implementing the described procedure could be implemented as a network App connected to the north-bound interface of an SDN controller. Such kinds of possible developments are in progress and will also be described in future papers.

Acknowledgement: The work of the Italian authors was partially sponsored by EU FP7 IRSES MobileCloud Project (Grant No. 612212) and by EIT ICT-Lab SOFNETS Project. M. Pióro was supported by the FP7 project IDEALIST (grant agreement no. 317999), and by National Science Centre, Poland (grant 2011/01/B/ST7/02967).

References

- [1] S.S. Ahuja, T. Korkmaz, and M. Krunz. Minimizing the differential delay for virtually concatenated Ethernet over SONET systems. In *23rd Int. Conf. Comput. Commun. and Networks (ICCCN)*, pages 205–210, Oct. 2004.
- [2] M. Alicherry, C. Phadke, and V. Poosala. Delay distributed VCAT for efficient data-optical transport. In *OFC*, volume 4, page 3 pp., Mar. 2005.
- [3] A. Anurag Srivastava, S. Acharya, M. Alicherry, B. Gupta, and P. Risbood. Differential delay aware routing for Ethernet over SONET/SDH. In *IEEE INFOCOM*, volume 2, pages 1117–1127, Mar. 2005.
- [4] M. Coudron, S. Secci, G. Pujolle, P. Raad, and P. Gallard. Cross-layer cooperation to boost multipath TCP performance in cloud networks. In *IEEE 2nd Int. Conf. on Cloud Networking (CloudNet)*, pages 58–66, Nov. 2013.
- [5] A. Ford, C. Raiciu, M. Handley, and Bonaventure O. TCP extensions for multipath operation with multiple addresses. RFC6824. Jan. 2013.
- [6] F. Hu, Q. Hao, and K. Bao. A survey on software-defined network and openflow: From concept to implementation. *Commun. Surveys Tuts.*, 16(4):2181–2206, Fourthquarter 2014.
- [7] H. Sheng, C.U. Martel, and B. Mukherjee. Survivable multipath provisioning with differential delay constraint in telecom mesh networks. *IEEE/ACM Trans. Netw.*, 19(3):657–669, June 2011.
- [8] A. Tomaszewski and M. Pióro. Minimizing label usage in MPLS networks. In *43rd Asilomar Conf. Signals, Syst. and Comput.*, pages 58–66, Nov. 2009.