



A Penalized Best-Response Algorithm for Non-Linear Single-Path Routing Problems

Olivier Brun, Balakrishna Prabhu, Josselin Vallet

► To cite this version:

Olivier Brun, Balakrishna Prabhu, Josselin Vallet. A Penalized Best-Response Algorithm for Non-Linear Single-Path Routing Problems. 2015. hal-01245236v1

HAL Id: hal-01245236

<https://hal.science/hal-01245236v1>

Preprint submitted on 16 Dec 2015 (v1), last revised 14 Sep 2016 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Penalized Best-Response Algorithm for Non-Linear Single-Path Routing Problems*

O. Brun, B.J. Prabhu, J. Vallet

CNRS, LAAS, 7 Av. du Colonel Roche, F-31400 Toulouse, France
Univ. Toulouse, LAAS, F-31400 Toulouse, France

October 5, 2015

Abstract

This paper is devoted to non-linear single path routing problems, which are known to be NP-hard even in the simplest cases. For solving these problems, we propose an algorithm inspired from Game Theory in which individual flows are allowed to independently select their path to minimize their own cost function. We design the cost function of the flows so that the resulting Nash equilibrium of the game provide an efficient approximation of the optimal solution. We establish the convergence of the algorithm and show that every optimal solution is a Nash equilibrium of the game. We also prove that if the objective function is a polynomial of degree $d \geq 1$, then the approximation factor of the algorithm is $(2^{1/d} - 1)^{-d}$. Experimental results show that the algorithm provides single-path routings with modest relative errors with respect to optimal solutions, while being several orders of magnitude faster than existing techniques.

1 Introduction

The single-path routing problem, also known as the unsplittable or non-bifurcated multi-commodity flow problem, naturally arises in a variety of contexts and, in particular, plays a central role in the traffic engineering of communication networks. It amounts to routing a given set of fixed traffic demands in a network, allocating each demand to a single path so as to minimize the total network cost, which is usually expressed as the sum of link costs (see [28, 19, 23] and Chapter 8 of [24] for the related joint routing and congestion control problem). A common assumption is that the cost of a link is a linear function of

*This is an extended version of the paper "A game-theoretic algorithm for non-linear single-path routing problems" that appeared in the proceedings of INOC 2015.

the traffic flowing on that link. Under this assumption, the single-path routing problem can be formulated as a mixed integer linear programming (MILP) problem with binary variables [2, 4, 5, 6]. Although this assumption might be appropriate in some contexts, it is no more viable when the cost of a link is interpreted as the delay incurred by a packet on that link since this delay depends non-linearly on the amount of traffic flowing over that link.

This paper is devoted to single-path routing problems with an additive and non-linear objective function. Our interest in these problems originates from the traffic engineering in communication networks based on the Multi-Protocol Label Switching (MPLS) technology [11]. MPLS changes the usual hop-by-hop paradigm by enabling network flows to be routed along predetermined paths, which are called Label Switched Paths (LSPs). In an MPLS network, each flow is routed along a single path, but two different flows with the same source/destination pair can use two different paths. Traffic engineering in MPLS networks mainly amounts to optimizing the quality of service of network flows by tailoring the paths assigned to flows to the prevailing traffic conditions. As explained above, when the quality of service metric is the network delay, the objective function depends non-linearly on the traffic flowing on each link and the problem at hand becomes a non-linear single path routing problem.

We aim to keep the discussion as general as possible by proposing a solution strategy that is also applicable in other application areas. We therefore focus on the essence of the non-linear single-path routing problem and omit on purpose many application-specific details of traffic engineering in MPLS networks. The problem we consider can be formulated as follows. Given some fixed traffic demands, the goal is to find a single-path routing strategy that minimizes the sum of link costs. The cost of a link e depends on the amount of traffic y_e flowing on that link (measured in packets/s) and has the form $y_e \ell_e(y_e)$, where the function ℓ_e gives the delay per packet on link e as a function of y_e . We specifically assume that the link latency function ℓ_e is a non-linear function. A typical example is the M/M/1 delay function $\ell_e(x) = 1/(c_e - x)$, where c_e represents the capacity of link e . Note however that the results established in this paper hold for a broad range of cost functions.

The routing problem as defined above belongs to the class of non-linear mathematical programs involving integer variables (actually, binary variables). These mathematical programs are known to be extremely hard to solve, both from a theoretical and from a practical point of view (see Chapter 15 of [14] for a survey). Even in the simplest case with binary variables, quadratic function and equality constraints, they are known to be NP-hard [9]. Several general approaches have been proposed to solve non-linear integer programming problems. Some transform the discrete problem into a continuous one (see for example [22]). Despite their qualities, those techniques do not scale very well with the size of the problems. A recent alternative is the so-called *Global Smoothing Algorithm* [21], which seems to scale better while providing fairly good approximations (see Section 5.1 for details). Heuristics and meta-heuristics have also been used to find an approximate solution to non-linear integer programming problems. Among others, ant-inspired opti-

mization techniques are known to be efficient for solving various routing problems [26], [18]. In this paper, we use the heuristic method proposed in [18] for comparison purposes.

We propose an approximation algorithm, inspired from Game Theory [17], for solving non-linear single-path routing problems. We shall assume that individual demands are allowed to independently select their path to minimize their own cost function. We design the cost function of the demands so that the resulting Nash equilibrium of the game provide an efficient approximation of the optimal solution. We note that a similar algorithm was proposed in [1] for scheduling of strictly periodic tasks. We establish the convergence of the algorithm and show that every optimal solution is a Nash equilibrium of the game. We also prove that if the link latency functions ℓ_e are polynomial of degree $d \geq 0$, then the approximation factor of the algorithm is $(2^{1/(d+1)} - 1)^{-(d+1)}$. As will be shown numerically, the main merit of this algorithm is that it is several orders of magnitude quicker than the method discussed above while providing good optimization results.

The rest of this paper is organized as follows. We introduce our notations and formally state the problem in Section 2. We then describe the proposed algorithm in Section 3. In Section 4, we show the convergence of the algorithm and obtain worst-case guarantees on the quality of the solutions it produces. Section 5 is devoted to the empirical performance evaluation of the algorithm. Finally, some conclusions are drawn in Section 6.

2 Problem Statement

Consider a network represented by a directed graph $G = (V, E)$. To each edge $e \in E$ is associated a nondecreasing latency function $\ell_e : \mathbb{R}_+ \rightarrow \mathbb{R}_+$. For any set $\pi \subset E$, we define the constant δ_π^e as 1 if $e \in \pi$, and 0 otherwise.

We are given a set $\mathcal{K} = \{1, 2, \dots, K\}$ of Origin-Destination (OD) pairs. We let s_k and t_k be the origin and destination of OD pair k , and $\lambda_k \in \mathbb{N}$ be its traffic demand. Each traffic demand has to be routed in the network over a single path. We let Π_k be the set of all paths available for routing traffic between s_k and t_k . This set can contain all simple paths from s_k to t_k , or only a subset of those paths satisfying some constraints. We define a routing strategy as a vector $\boldsymbol{\pi} = (\pi_k)_{k \in \mathcal{K}} \in \Pi$, where π_k is the path assigned to traffic demand k and $\Pi = \Pi_1 \times \Pi_2 \times \dots \times \Pi_K$. The goal is to find a routing strategy that minimizes the cost of the network $F(\boldsymbol{\pi}) = \sum_{e \in E} y_e(\boldsymbol{\pi}) \ell_e(y_e(\boldsymbol{\pi}))$, where $y_e(\boldsymbol{\pi}) = \sum_{k \in \mathcal{K}} \delta_{\pi_k}^e \lambda_k$ is the total traffic flowing on link e in routing strategy $\boldsymbol{\pi}$. Formally, the problem is as follows:

$$\text{minimize } F(\boldsymbol{\pi}) = \sum_{e \in E} y_e(\boldsymbol{\pi}) \ell_e(y_e(\boldsymbol{\pi})) \quad (\text{OPT})$$

subject to

$$\pi_k \in \Pi_k \quad k \in \mathcal{K}. \quad (1)$$

We note that the problem can be formulated as a 0-1 mathematical programming problem by introducing the binary variables

$$x_{k,\pi} = \begin{cases} 1 & \text{if demand } k \text{ is routed on path } \pi \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

Indeed, we can rewrite problem (OPT) as follows:

$$\begin{aligned} & \text{minimize } \sum_{e \in E} y_e \ell_e(y_e) \\ & \text{subject to} \end{aligned}$$

$$y_e = \sum_{k \in \mathcal{K}} \sum_{\pi \in \Pi_k} \lambda_k \delta_{\pi}^e x_{k,\pi} \quad e \in E, \quad (3)$$

$$\sum_{\pi \in \Pi_k} x_{k,\pi} = 1 \quad k \in \mathcal{K}, \quad (4)$$

$$x_{k,\pi} \in \{0, 1\} \quad \pi \in \Pi_k, k \in \mathcal{K}. \quad (5)$$

When the objective function F is non-linear function of the variables y_e , the above problem belongs to the class of non-linear mathematical programs involving integer variables (actually, binary variables in our case). These mathematical programs are known to be extremely hard to solve, thus motivating the development of efficient methods for finding approximate solutions.

3 Best-response algorithm

The best-response algorithm we propose takes its inspiration from an algorithm of the same name in Game theory. In a game, the best-response of a player is defined as its optimal strategy conditioned on the strategies of the other players. It is, as the name suggests, the best response that the player can give for a given strategy of the others. The best-response algorithm then consists of players taking turns in some order to adapt their strategy based on the most recent known strategy of the others until an equilibrium point is reached.

In this section, we design the game so that the Nash equilibrium of the game provides an efficient approximation to problem (OPT). We follow the approach proposed in [10]. Let us think of the traffic demands as the players of the game. The strategy of player k is the path π_k it chooses in the set Π_k , and a strategy profile of the game is a vector $\boldsymbol{\pi} \in \Pi$. In other words, a strategy profile corresponds to a feasible solution of problem (OPT). Given the strategy of the other players $\boldsymbol{\pi}_{-i} = (\pi_1, \pi_2, \dots, \pi_{i-1}, \pi_{i+1}, \dots, \pi_K)$, the value $f_i(\pi, \boldsymbol{\pi}_{-i})$ associated to path $\pi \in \Pi_i$ by player i reflects the cost of this path, i.e.,

$$f_i(\pi, \pi_{-i}) = \sum_{e \in \pi} \lambda_i \ell_e(y_e(\pi, \pi_{-i})), \quad \forall \pi \in \Pi_i \quad (6)$$

We note that

$$F(\pi) = \sum_e \left(\sum_k \delta_{\pi_k}^e \lambda_k \right) \ell_e(y_e(\pi)) = \sum_k \lambda_k \sum_{e \in \pi_k} \ell_e(y_e(\pi)) = \sum_k f_k(\pi). \quad (7)$$

When player i minimizes the cost given in (6), this game is called a weighted congestion game (see [7] and references therein). For these games, several authors, including [7], have investigated what is called the *Price of Anarchy* which is the ratio of the social cost (the function F in our case) at worst-case Nash equilibrium and the social cost at the global optimum. Although one could design a best-response algorithm based on (6), that algorithm will have two drawbacks: (i) a global optimum of F need not be a Nash equilibrium of the weighted congestion game as is shown in example 1; and (ii) the convergence of the best-response algorithm for the weighted congestion game is known only in some special cases (linear latency functions or when the demands are the same).

Example 1. Consider a network with two parallel links as shown in figure 1. There are two flows, each of size 1, which wish to go from O to D . There are two possible routes: (i) the top-route with a latency function of $l(x) = x$; and (ii) the bottom-route with a latency function of $l(x) = 0.4x$.




Figure 1: Example to show that the Nash equilibrium of a weighted congestion game need not be a global optimum.

It can be verified that at the global optimum one of the flows will be routed through the top-route and the other will be routed through the bottom-route. The cost for the flow on the top-route at the optimal routing is 1 while for the one on the bottom-route it is 0.4. In order to check that the optimum is not a Nash equilibrium for cost defined in (6), we look at the best-response of the flow on the top-route. If this flow moves to the bottom-route its cost will be $0.4(1 + 1) = 0.8$ which is less than its current cost of 1. Thus, at the global optimum the flow on the top-route has an incentive to move to the bottom-route.

In order to drive the players towards a local minimum of $F(\boldsymbol{\pi})$, we add a penalty term $p_i(\pi, \boldsymbol{\pi}_{-i})$ to the cost of player i . The role of the penalty term is that of an incentive for players to take into account the impact of their actions on other players. When player i routes its traffic on path π , the increase in the cost of each player $j \neq i$ using link $e \in \pi$ is

$$\lambda_j [\ell_e (y_e^{-i} + \lambda_i) - \ell_e (y_e^{-i})], \quad (8)$$

where $y_e^{-i} = \sum_{k \neq i} \delta_{\pi_k}^e \lambda_k$ represents the total traffic flowing on link e due to all players other than i . As a consequence, we define the penalty term as follows:

$$\begin{aligned} p_i(\pi, \boldsymbol{\pi}_{-i}) &= \sum_{e \in \pi} \sum_{j \neq i} \lambda_j \delta_{\pi_j}^e [\ell_e (y_e^{-i} + \lambda_i) - \ell_e (y_e^{-i})], \\ &= \sum_{e \in \pi} y_e^{-i} [\ell_e (y_e^{-i} + \lambda_i) - \ell_e (y_e^{-i})]. \end{aligned} \quad (9)$$

In summary, player i computes its path so as to minimize its own cost, that is, player i solves the following problem:

$$\text{minimize}_{\pi \in \Pi_i} c_i(\pi, \boldsymbol{\pi}_{-i}) = f_i(\pi, \boldsymbol{\pi}_{-i}) + p_i(\pi, \boldsymbol{\pi}_{-i}). \quad (\text{OPT-}i)$$

In (OPT- i), the path π minimizing $c_i(\pi, \boldsymbol{\pi}_{-i})$ is known as the best response of player i to the strategy $\boldsymbol{\pi}_{-i}$ of the other players. The game is in a Nash equilibrium if and only if the current strategy of each player is its best response to the strategy of the others, implying that no player has an incentive to unilaterally deviate from its current strategy. Formally, $\boldsymbol{\pi}$ is a Nash equilibrium if and only if

$$c_i(\boldsymbol{\pi}) \leq c_i(\pi, \boldsymbol{\pi}_{-i}), \forall \pi \in \Pi_i, \forall i \in \mathcal{K}.$$

As we will show briefly, the best-response algorithm that we propose as an approximation to problem (OPT) computes a Nash equilibrium. The algorithm starts from an initial feasible solution $\boldsymbol{\pi}^{(0)}$. At each iteration, the players update their strategies in a given order by computing the optimal solution of problem (OPT- i) using any shortest path algorithm (e.g., Dijkstra's algorithm). Note that a player i deviates from its strategy $\pi_i^{(n)}$ at iteration n to a new strategy π' if and only if $c_i(\pi', \boldsymbol{\pi}_{-i}^{(n)}) < c_i(\pi^{(n)})$. The algorithm stops when no player can decrease its cost by unilaterally deviating from its strategy, that is $\boldsymbol{\pi}^{(n+1)} = \boldsymbol{\pi}^{(n)}$. The pseudocode for the heuristic is given in Algorithm 1.

In the following, we shall use the term *penalized best-response* to refer to Algorithm 1. The term *standard best-response algorithm* will refer to the case when each player i directly optimizes $f_i(\pi, \boldsymbol{\pi}_{-i})$ instead of the penalized cost function $c_i(\pi, \boldsymbol{\pi}_{-i}) = f_i(\pi, \boldsymbol{\pi}_{-i}) + p_i(\pi, \boldsymbol{\pi}_{-i})$.

Algorithm 1 Penalized best-response

Require: $\pi^{(0)}$

```
1:  $n \leftarrow 0$ 
2: repeat
3:   for  $i = 1 \dots K$  do
4:      $\pi_i^{(n+1)} \leftarrow \text{argmin}(\text{OPT-}i)$ 
5:   end for
6:    $n \leftarrow n + 1$ 
7: until  $\pi^{(n+1)} \neq \pi^{(n)}$ .
8: return  $\vec{\pi}(n)$ 
```

4 Properties of the penalized best-response algorithm

In this section, we establish several properties of the penalized best-response algorithm. We first show in Section 4.1 that the penalized best-response algorithm converges in a finite number of steps to a Nash equilibrium. In contrast, the convergence of the standard best-response algorithm is known only (a) for linear latency functions [16], and (b) when all flows have the same traffic demands (that is when $\lambda_k = \lambda, \forall k$), in which case the game is a congestion game [25].

In Section 4.2, we obtain performance guarantees for the solutions provided by the penalized best-response algorithm in the special case of polynomial link latency functions. In this special case, the algorithm is therefore an approximation algorithm since it returns a solution that is provably close to optimal (in contrast to other cases where it is only a heuristic that may or may not find a good solution). Our performance guarantees take the form of worst-case bounds on the approximation factor of our algorithm, that is bounds on the ratio between the result obtained by the algorithm and the optimal cost. We use techniques similar to those used in [3], where the authors obtain bounds on the approximation factor of the standard best-response algorithm for affine and polynomial cost functions with non negative coefficients.

4.1 Convergence of the penalized best-response algorithm

As proven in [10], the convergence of the penalized best-response algorithm directly follows from the fact that F is a potential function of the game. We briefly describe the proof below. The crucial argument of this proof is stated in the following lemma.

Lemma 1. *The objective function of player i can be written as*

$$c_i(\pi) = F(\pi) - h(\pi_{-i}), \quad (10)$$

where

$$h(\boldsymbol{\pi}_{-i}) = \sum_{e \in E} y_e^{-i} \ell_e(y_e^{-i}), \quad (11)$$

does not depend on the path π_i chosen by player i .

Proof. Note that $y_e(\boldsymbol{\pi}) = y_e^{-i}$ for $e \notin \pi_i$, while $y_e(\boldsymbol{\pi}) = y_e^{-i} + \lambda_i$ for $e \in \pi_i$. Hence

$$\begin{aligned} c_i(\boldsymbol{\pi}) &= f_i(\pi_i, \boldsymbol{\pi}_{-i}) + p_i(\pi_i, \boldsymbol{\pi}_{-i}), \\ &= \sum_{e \in \pi_i} \lambda_i \ell_e(y_e^{-i} + \lambda_i) + \sum_{e \in \pi_i} y_e^{-i} [\ell_e(y_e^{-i} + \lambda_i) - \ell_e(y_e^{-i})], \\ &= \sum_{e \in \pi_i} (y_e^{-i} + \lambda_i) \ell_e(y_e^{-i} + \lambda_i) - \sum_{e \in \pi_i} y_e^{-i} \ell_e(y_e^{-i}), \\ &= \left(\sum_{e \in \pi_i} y_e(\boldsymbol{\pi}) \ell_e(y_e(\boldsymbol{\pi})) + \sum_{e \notin \pi_i} y_e(\boldsymbol{\pi}) \ell_e(y_e(\boldsymbol{\pi})) \right) - \left(\sum_{e \in \pi_i} y_e^{-i} \ell_e(y_e^{-i}) + \sum_{e \notin \pi_i} y_e^{-i} \ell_e(y_e^{-i}) \right), \\ &= F(\boldsymbol{\pi}) - h(\boldsymbol{\pi}_{-i}). \end{aligned}$$

□

Proposition 1. *The penalized best-response algorithm converges in a finite number of steps to a Nash equilibrium.*

Proof. A direct consequence of lemma 1 is that the function F is a potential function of the game, i.e.,

$$c_i(\boldsymbol{\pi}) - c_i(\pi, \boldsymbol{\pi}_{-i}) = F(\boldsymbol{\pi}) - F(\pi, \boldsymbol{\pi}_{-i}), \quad \forall \pi \in \Pi_i, \forall i \in \mathcal{K},$$

implying that any best-response move results in a decrease of the global cost $F(\boldsymbol{\pi})$. Since $F(\boldsymbol{\pi})$ can accept a finite number of values, the sequence will reach a local minimum in a finite number of steps. □

Another consequence of lemma 1 is that any global optimum is a Nash equilibrium, so that the penalized best-response algorithm stops in an optimal point if it ever reaches it.

Proposition 2. *Any global optimum is a Nash equilibrium of the routing game.*

Proof. Consider a global optimum $\boldsymbol{\pi}^*$, i.e., a point such that $F(\boldsymbol{\pi}^*) \leq F(\boldsymbol{\pi})$ for all $\boldsymbol{\pi} \in \times_i \Pi_i$. Assume on the contrary that $\boldsymbol{\pi}^*$ is not a Nash equilibrium. It follows that there exists a player i and a strategy $\pi_i \in \Pi_i$ such that $c_i(\pi_i, \boldsymbol{\pi}_{-i}^*) < c_i(\boldsymbol{\pi}^*)$. With lemma 1, it yields $F(\pi_i, \boldsymbol{\pi}_{-i}^*) - h(\boldsymbol{\pi}_{-i}^*) < F(\boldsymbol{\pi}^*) - h(\boldsymbol{\pi}_{-i}^*)$, from which we conclude that $F(\pi_i, \boldsymbol{\pi}_{-i}^*) < F(\boldsymbol{\pi}^*)$. This is clearly a contradiction with the global optimality of $\boldsymbol{\pi}^*$. Hence $\boldsymbol{\pi}^*$ is a Nash equilibrium. □

4.2 Performance guarantees for polynomial link latency functions

An instance I of our problem is defined by the graph $G = (V, E)$, by the set of traffic demands \mathcal{K} (including the demand value λ_k and the set of candidate paths Π_k for routing each demand k) and by the link latency functions $\ell_e : \mathbb{R}_+ \rightarrow \mathbb{R}_+$. Our goal in this section, is to establish bounds on the approximation ratio of the penalized best-response algorithm that hold uniformly over all instances of the problem. More precisely, given an instance I of the problem, let π be any Nash equilibrium and π^* be an optimal routing strategy for that instance. We look for an upper bound on the ratio $F(\pi)/F(\pi^*)$ that hold for all instances I of the problem. In the following, we establish such a bound when the link latency functions are polynomial functions of the form $\ell_e(x) = \sum_{j=0}^d a_{e,j} x^j$, where $d \geq 0$ and the a_e are non-negative coefficients. Our main result is stated in theorem 3.

Theorem 3. *If $\ell_e(x) = \sum_{j=0}^d a_{e,j} x^j$, then the approximation factor of the penalized best-response algorithm is upper bounded by $(2^{1/(d+1)} - 1)^{-(d+1)}$.*

Proof. See appendix A. □

We can furthermore prove that the upper bound of theorem 3 is tight, in the sense that there exists an instance of the single-path routing problem for which it is reached. The lower bound on the approximation ratio is stated in proposition 4, whose proof is based on an adaptation of the instance given in [7] (see Lower Bound 2 and lemma 4.5 in that paper).

Proposition 4. *If $\ell_e(x) = \sum_{j=0}^d a_{e,j} x^j$, then the approximation ratio of the penalized best-response algorithm is lower bounded by $(2^{1/(d+1)} - 1)^{-(d+1)}$.*

Proof. We consider a problem instance with a monomial cost function, that is, $\ell_e(x) = a_e x^d$, and N origin-destination pairs. In addition to these nodes, there are two other nodes R_0 and R_1 which act as routers. Origin O_i is connected to nodes O_{i-1} and O_{i+1} , except O_1 which is connected to R_0 and O_2 , and O_N which is connected to O_{N-1} and R_1 . Destination D_i is connected to O_{i-1} and O_{i+1} , except D_1 which is connected to R_0 and O_2 , and D_N which is connected to O_{N-1} and R_1 . In the following, we shall assume that links are numbered in such a way that link 0 is link $R_0 O_1$, link i is link $O_i O_{i+1}$ for $i = 1, \dots, N-1$ and link N is link $O_N R_1$. For $i = 0, \dots, N$, link i has a cost of $a_i y^{d+1}$ when it carries a traffic of y , whereas for all other links we have $a_e = 0$. The instance has a total of $3N + 1$ links of which $2N$ links have a cost of 0 and the $N + 1$ other links have a non-zero cost. Figure 2 shows an instance with $N = 4$ pairs.

Let λ_i be the amount of traffic sent from origin O_i to destination D_i . Each origin has two choices: either send the traffic to the left or to the right. In the global optimum (OPT) each origin will send its traffic to the right provided that

$$a_{i+1} \leq a_i \quad \text{and} \quad a_N = a_{N_1}. \quad (12)$$




Figure 2: Problem instance with $N = 4$ origin-destination pairs for the lower bound on the approximation ratio.

Indeed, in that case, each flow is routed over the link with the smaller a_e , and over another link of cost $\ell_e(x) = 0$. Similarly, we shall establish the conditions under which routing the traffic to the left corresponds to a Nash equilibrium (NE). It is enough, according to lemma 1, to show that we cannot decrease the network cost by deviating a single flow from its equilibrium route onto another route. Provided that $a_N = a_{N-1}$, the cost does not change when flow N is routed to the right on link $O_N R_1$ instead of being routed to the left on link $O_{N-1} O_N$. If flow $i = 1, \dots, N_1$ is routed to the right on link i instead of being routed to the left on link $i - 1$, the traffic on link i increases from λ_{i+1} to $\lambda_i + \lambda_{i+1}$, whereas the traffic on link i becomes 0. We deduce that the cost variation is

$$a_i(\lambda_i + \lambda_{i+1})^{d+1} + 0 - \left[a_{i-1}\lambda_i^{d+1} + a_i\lambda_{i+1}^{d+1} \right],$$

so that routing to the left corresponds to a Nash equilibrium provided that

$$a_{i-1}\lambda_i^{d+1} + a_i\lambda_{i+1}^{d+1} \leq a_i(\lambda_i + \lambda_{i+1})^{d+1}, \quad i = 1, \dots, N-1, \quad (13)$$

and $a_N = a_{N-1}$. Let us now assume that $a_i = \phi^{(d+1)i}$ for $0 \leq i < N$ and that $a_N = \phi^{(d+1)(N-1)}y^{d+1}$, where $\phi = 2^{1/(d+1)} - 1$. Note that the links $O_{N-1}O_N$ and $O_N R_1$ have the same cost. In figure 2, for the links with a non-zero cost, the corresponding a_e is shown close to the mid-point of the link. We further assume that $\lambda_i = \phi^{N-i}$. Since $\phi < 1$, $a_N = a_{N-1}$ and $a_{i+1} = \phi^{(d+1)}a_i$ for $i = 1, \dots, N-1$ implies that the conditions (12) are satisfied, so that routing to the right is indeed an optimal strategy. To see that routing to the left corresponds to a Nash equilibrium, observe that for $i < N$,

$$\begin{aligned} a_{i-1}\lambda_i^{d+1} + a_i\lambda_{i+1}^{d+1} &= 2\phi^{(d+1)(N-1)} \\ &= (1 + \phi)^{d+1} \phi^{(d+1)(N-1)} \\ &= a_i(\lambda_i + \lambda_{i+1})^{d+1}, \end{aligned}$$

Table 1: Traffic on the links and associated cost in the optimal solution and at the Nash equilibrium for the instance shown in figure 2.

	NE			OPT		
Links	Flow	Traffic	Cost	Flow	Traffic	Cost
R_0O_1	1	ϕ^3	$\phi^{3(d+1)}$	—	0	0
O_1O_2	2	ϕ^2	$\phi^{3(d+1)}$	1	ϕ^3	$\phi^{4(d+1)}$
O_2O_3	3	ϕ^1	$\phi^{3(d+1)}$	2	ϕ^2	$\phi^{4(d+1)}$
O_3O_4	4	ϕ^0	$\phi^{3(d+1)}$	3	ϕ^1	$\phi^{4(d+1)}$
O_4R_1	—	0	0	4	ϕ^0	$\phi^{3(d+1)}$

which proves that the network cost cannot be decreased by deviating unilaterally any flow. It is interesting to note that this instance has been designed in such a way that the corresponding Nash equilibrium is a weak equilibrium, that is, for each flow i , there exists a path π different from its equilibrium path π_i such that $c_i(\pi, \pi_{-i}) = c_i(\pi)$.

For the instance with $N = 4$, in table 1, we give the traffic on each link and the associated cost for each of the two solutions (for the clarity of the presentation, only links e such that $\ell_e(x) \neq 0$ are presented). Summing the costs of the links in table 1, we observe that the optimal cost is $F_4(\pi^*) = 3\phi^{4(d+1)} + \phi^{3(d+1)}$, whereas the cost at the Nash equilibrium is $F_4(\pi) = 4\phi^{3(d+1)}$.

In general, it can be seen that $F_N(\pi^*) = (N-1)\phi^{(d+1)N} + \phi^{(d+1)(N-1)}$ and that $F_N(\pi) = N\phi^{(d+1)(N-1)}$. As a consequence, the approximation ratio is

$$g_N(\phi) = \frac{N}{(N-1)\phi^{d+1} + 1}. \quad (14)$$

By taking $N \rightarrow \infty$, we obtain an approximation ratio of $\phi^{-(d+1)} = (2^{1/(d+1)} - 1)^{-(d+1)}$ as claimed. \square

Remark 1. *For the sake of clarity, we have shown the lower bound on an undirected graph. By adding links and intermediate nodes, this instance can be modified to make the graph directed as well. However, this will make the figure and the notation cumbersome.*

As an immediate consequence of theorem 3 and proposition 4, we obtain the following corollary.

Corollary 5. *For large d , the approximation factor is roughly $\left(\frac{d+1}{\log(2)}\right)^{d+1}$.*

In table 2, we give the approximation factors for different values of d for the penalized best-response algorithm and compare it with that of the non-penalized version, as established in [3] and [7].

Table 2: Approximation factor for the penalized best-response as well as for the non-penalized best-response as a function of the degree of the polynomial. The values have been rounded-off at two decimal places.

degree (d)	Penalized BR	Standard BR [7]
0	1	1
1	5.83	2.62
2	56.95	9.91
3	780.28	47.82
asymptotic	$\Theta\left(\left(\frac{d+1}{\log(2)}\right)^{d+1}\right)$	$\Theta\left(\left(\frac{d}{\log(d)}\right)^{d+1}\right)$

Interestingly, in the case $d = 0$, that is when the latency function $\ell_e(x)$ of each link e is a constant $a_e \geq 0$, our algorithm is guaranteed to provide an optimal solution. This is a direct consequence of theorem 3, but can also be easily understood by noting that in this case

$$F(\boldsymbol{\pi}) = \sum_{e \in E} a_e y_e(\boldsymbol{\pi}) = \sum_{e \in E} a_e \sum_{k \in \mathcal{K}} \delta_{\pi_k}^e \lambda_k = \sum_{k \in \mathcal{K}} \lambda_k \left(\sum_{e \in \pi_k} a_e \right),$$

so that it is optimal to route each flow k on the path $\pi_k \in \Pi_k$ minimizing $\sum_{e \in \pi_k} a_e$, which can be seen as the length of the path. However, in that case (9) yields $p_i(\boldsymbol{\pi}) = 0$, from which it follows that $c_i(\boldsymbol{\pi}, \boldsymbol{\pi}_{-i}) = f_i(\boldsymbol{\pi}, \boldsymbol{\pi}_{-i}) = \lambda_i (\sum_{e \in \pi_i} a_e)$, so that our algorithm routes each flow on a shortest path, exactly as does the optimal solution.

When $d \geq 1$, the approximation ratio of the penalized best-response algorithm is significantly worse than that of the non-penalized best-response algorithm. Asymptotically, the non-penalized version is a factor $(\log(d))^{d+1}$ better than the penalized version. On the downside, the non-penalized version does not have a guaranteed convergence whereas the penalized one is guaranteed to converge.

It is however important to keep in mind that the above performance guarantees are obtained for worst-case scenarios that are not necessarily representative of instances met in practice. We conjecture that the worst-case performance of the algorithm occur in very asymmetric scenarios in terms of link latency functions and in terms of traffic demands. Consider for example the instance used in proposition 4 to obtain a lower bound on the approximation ratio. Assume that the link latency functions are monomials of degree d such that $a_N = a_{N-1}$ and $1 \leq \frac{a_i}{a_{i+1}} \leq x$ for some fixed x such that

$$1 \leq x \leq \min_i \frac{(\lambda_i + \lambda_{i+1})^{(d+1)} - \lambda_{i+1}^{(d+1)}}{\lambda_i^{d+1}}.$$

The above conditions impose a certain homogeneity between the coefficients a_i . It is readily verified that equations (12) and (13) are satisfied under these conditions, so that routing to the right is an optimal strategy, and routing to the left is a Nash equilibrium. Observe now that

$$\frac{F(\boldsymbol{\pi})}{F(\boldsymbol{\pi}^*)} = \frac{\sum_{i=0}^{N-1} a_i \lambda_{i+1}^{(d+1)}}{\sum_{i=1}^N a_i \lambda_i^{(d+1)}} \leq x \frac{\sum_{i=0}^{N-1} a_{i+1} \lambda_{i+1}^{(d+1)}}{\sum_{i=1}^N a_i \lambda_i^{(d+1)}} = x,$$

so that in this example we can obtain a solution as close as we want to an optimal solution by imposing a condition on the ratio $\frac{a_i}{a_{i+1}}$. Unfortunately, we were not able to explicitly characterize the conditions under which the approximation ratio will be lower than a given x in the general case. Nevertheless, the numerical experiments presented in Section 5 reveal that the penalized best-response often provides good quality solutions, contrary to what is suggested by the approximation ratio $(2^{1/(d+1)} - 1)^{-(d+1)}$. These experiments also show that the results of the penalized best-response algorithm are quite similar to those of its non-penalized counter-part.

5 Empirical Performance Evaluation

This section is devoted to the empirical performance evaluation of the best-response algorithm. We first describe the benchmark experiments that were performed, as well as our empirical algorithm comparison methodology in Section 5.1. Section 5.2 reports the numerical results obtained.

5.1 Benchmark Experiments

In our benchmark experiments, we have used 800 randomly generated instances obtained using 8 standard network topologies (see Table 3) collected from the IEEE literature and from the Rocketfuel project [27]. For each network topology, we consider 100 random traffic matrices generated with uniform distributions in such a way that there is positive traffic demand associated to each origin/destination pair. The traffic matrices are computed so that there exists a minimum-hop routing strategy in which the network congestion rate (that is, the maximum utilization rate of the network links) is equal to γ , where $\gamma > 0$ is a given parameter.

The empirical evaluation of the best-response algorithm is done for two different types of objective functions, as described below.

5.1.1 Piece-wise Linear Objective Function

The main motivation for considering a piece-wise linear objective function is that in that case we can compare the algorithm with the optimal solutions obtained from a MILP solver. Specifically, we consider the following increasing latency function proposed in [15]:

Table 3: Topologies : number of nodes and links.

Topology	# nodes	# links
ABOVENET	19	68
ARPANET	24	100
BHVAC	19	46
EON	19	74
METRO	11	84
NSF	8	20
PACBELL	15	42
VNSL	9	22

$$\ell_e(y) = \begin{cases} 1 & \text{if } 0 \leq \frac{y}{c_e} < \frac{1}{3}, \\ 3 - \frac{2}{3} \frac{c_e}{y} & \text{if } \frac{1}{3} \leq \frac{y}{c_e} < \frac{2}{3}, \\ 10 - \frac{16}{3} \frac{c_e}{y} & \text{if } \frac{2}{3} \leq \frac{y}{c_e} < \frac{9}{10}, \\ 70 - \frac{178}{3} \frac{c_e}{y} & \text{if } \frac{9}{10} \leq \frac{y}{c_e} < 1, \\ 500 - \frac{1468}{3} \frac{c_e}{y} & \text{if } 1 \leq \frac{y}{c_e} < \frac{11}{10}, \\ 5000 - \frac{16318}{3} \frac{c_e}{y} & \text{if } \frac{11}{10} \leq \frac{y}{c_e}, \end{cases}$$

where c_e represents the capacity of link e . This function expresses that it is cheap to send flow over a link with a small utilization rate, whereas, as the load approaches 100%, it becomes more expensive. Note that the utilization rate can be greater than 1, in which case we get heavily penalized. For the above latency function, we assume that the set Π_k contains all the possible path between s_k and t_k for all flows k . Problem (OPT) can then be formulated as a Mixed-Integer Linear Problem:

$$\begin{aligned}
& \text{minimize } \sum_{e \in E} \Phi_e && \text{(OPT-PWL)} \\
& \text{subject to} \\
& \sum_{e \in In(n)} x_{k,e} - \sum_{e \in Out(n)} x_{k,e} = \begin{cases} -1 & \text{if } n = s, \\ 1 & \text{if } n = t, \\ 0 & \text{otherwise.} \end{cases} && n \in V, k \in \mathcal{K}, \quad (15) \\
& y_e = \sum_{k \in \mathcal{K}} \lambda_k x_{k,e} && e \in E, \quad (16) \\
& \Phi_e \geq ay_e - bc_e && (a, b) \in \mathcal{C}, e \in E, \quad (17) \\
& x_{k,e} \in \{0, 1\}, y_e \geq 0, \Phi_e \geq 0 && e \in E, k \in \mathcal{K}, \quad (18)
\end{aligned}$$

where $\mathcal{C} = \{(1, 0), (3, \frac{2}{3}), (10, \frac{16}{3}), (70, \frac{178}{3}), (500, \frac{1468}{3}), (5000, \frac{16318}{3})\}$ and $In(n)$ (resp. $Out(n)$) represents the set of all ingoing (resp. outgoing) links at node n . In the above node/link formulation, the decision variable $x_{k,e}$ is 1 if flow k is routed over link e , and 0 otherwise.

For each of the 800 instances, we have computed the relative error of the best-response algorithm with respect to the optimal solution of (OPT-PWL), which was computed using Gurobi 6.0 as MILP solver. We have limited the total time expended by the solver to 15 minutes per problem instance. The traffic matrices were computed with $\gamma = 1.2$ as threshold parameter, so that we can guarantee that there exists a feasible solution in which the utilization rate of the links is at most 120%.

5.1.2 Non-linear Objective Function

We consider two different link latency functions, each one yielding a non-linear network cost function. The first one is a linear latency function, $\ell_e(y) = \frac{y}{c_e}$, where c_e represents the capacity of link e , yielding the following quadratic objective function

$$F(\boldsymbol{\pi}) = \sum_{e \in E} \left(\frac{y_e(\boldsymbol{\pi})}{c_e} \right)^2. \quad (19)$$

The second link latency function corresponds to the well-known M/M/1 delay function (a.k.a. Kleinrock function) [20], that is $\ell_e(y) = \frac{1}{c_e - y}$, so that

$$F(\boldsymbol{\pi}) = \sum_{e \in E} \frac{y_e(\boldsymbol{\pi})}{c_e - y_e(\boldsymbol{\pi})}. \quad (20)$$

This cost function is often used when one wants to minimize the mean packet delay in the network. For the above two objective functions, it becomes extremely hard to find an optimal solution to problem (OPT). We therefore compare our penalized best-response

algorithm with a lower bound on the optimal cost obtained by relaxing the 0-1 constraints in problem (OPT), thereby solving the following multi-path routing problem

$$\begin{aligned} & \text{minimize} \sum_{e \in E} y_e \ell_e(y_e) & (\text{OPT-NL-MP}) \\ & \text{subject to} \end{aligned}$$

$$y_e = \sum_{k \in \mathcal{K}} \sum_{\pi \in \Pi_k} \lambda_k \delta_{\pi}^e x_{k,\pi} \quad e \in E, \quad (21)$$

$$\sum_{\pi \in \Pi_k} x_{k,\pi} = 1 \quad k \in \mathcal{K}, \quad (22)$$

$$0 \leq x_{k,\pi} \leq 1 \quad \pi \in \Pi_k, k \in \mathcal{K}, \quad (23)$$

which can easily be done with a projected gradient algorithm.

In addition to the comparison with the optimal multi-path solution, we compare our algorithm with two other algorithms designed to solve problem (OPT):

- **Global Smoothing Algorithm (GSA):** This algorithm was introduced in [21] for solving non-linear optimization problems with binary variables. Its original approach is to solve a sequence of non-linear optimization problems without integer variables. In each problem of the sequence, the goal is to minimize the sum of the original objective function plus two penalty terms: a logarithmic barrier penalty (designed to smooth the function and keep the current point away from the integer grid) and an “exact” penalty (describing how far the current point is from the integer grid). Both terms are weighted by one parameter, which is updated at each iteration in order to gradually bring the current point onto the integer grid. At each step, a continuous optimization problem is solved using a gradient algorithm (we have used the modified conjugate gradient algorithm presented in [8]). The sequence of penalized problem should form a trajectory leading to a good approximation of the solution. We refer to [21] for further details.
- **Ant Colony Optimization (ACO):** This heuristic algorithm was proposed in [18]. It belongs to the family of ant colony algorithms, which are known to be efficient for searching an optimal path in a graph [12, 26, 13]. The algorithm iteratively improves an initial solution to the problem through a random exploration of the solution space by agents called (artificial) ants. During one iteration of the algorithm, each ant builds a solution to the problem by randomly taking decisions, that is by randomly assigning paths to flows. A weight, called pheromone trail, is assigned to each decision taken by ants and is updated based on the experience acquired by ants during problem solving. These weights, which represent the learned desirability of the decisions, are used to guide the random exploration of the solution space.

Given the size and complexity of the problem instances, and in order to keep the running times of these algorithms below a reasonable amount, we restrict ourselves to two possible paths for each OD flow (that is, $|\Pi_k| = 2$ for all $k \in \mathcal{K}$). These paths are obtained by solving a 2-shortest-path problem in which the link weights are equal to 1. The traffic matrices were computed with $\gamma = 1.0$ as threshold parameter, so that we can guarantee that there exists a feasible solution in which the capacity of no link is exceeded, which is particularly important for the M/M/1 latency function.

The GSA and ACO algorithms have been implemented in Matlab 2013b. For GSA, two different initial multi-path solutions are considered: the first one is randomly generated, whereas the second one is obtained by adding a small random perturbation to the single-path solution obtained with the best-response algorithm (so as to obtain a multi-path initial solution). The results presented in the following for the GSA algorithms correspond to the best solution obtained from these two initial points. Regarding ACO, we consider at most 50 iterations and 5 ants since our experiments revealed that increasing these values does not improve significantly obtained results while negatively impacting computing times.

5.2 Numerical Results

5.2.1 Piece-wise Linear network cost function

The numerical results obtained for this cost function are reported in Table 4. Interestingly, the worst results are obtained for the smallest topology, NSF and METRO. Except for these topologies, the best-response algorithm performs quite well. Its average error over the 800 instances is only 3.31%, and it would be as low as 1.77% without the instances generated with the NSF topology. The average execution time of the algorithm is 0.51 seconds, and it never exceeds 3.2 seconds. This is compared to the average execution time of 82.6 seconds of the MILP solver (the time limit of 15 minutes was reached for 46 out of the 100 instances generated with the EON topology).

5.2.2 Quadratic network cost function

The numerical results obtained with this cost function are presented in Table 5. First of all, note that all three algorithms perform quite well since the average relative error is below 4.37% for each one. Over the 800 instances, the maximum relative error of the best-response algorithm is 20.16% (to be compared to the 13.32% of the other algorithms). We also note that on the average GSA performs slightly better than ACO and the best-response algorithm. Remember however that GSA is initialized with the solution obtained with the best-response algorithm. If we now look at the computing times presented in Table 7, we observe that the best-response algorithm is significantly faster than the two other algorithms, since its worst computing time is 6.8 seconds, whereas the running times of ACO and GSA can be as high as 50 and 300 seconds, respectively.

Table 4: Relative gap (%) to the optimal solution for the piece-wise linear objective function.

Topology	min	max	avg
ABOVENET	0.04	1.912	0.725
ARPANET	0.08	3.488	1.103
BHVAC	0.347	4.225	1.838
EON	0.938	6.390	3.022
METRO	0.177	10.297	3.969
NSF	0.081	63.021	14.128
PACKBEL	0.313	3.294	1.520
VNSL	0.010	2.556	0.240
<i>Global</i>	0.010	63.021	3.31

Table 5: Relative gap (%) to the optimal multi-path solution for the quadratic cost function.

Topology	BR			GSA			ACO		
	min	max	avg	min	max	avg	min	max	avg
ABOVENET	1,28	5,18	3,86	$\simeq 0$	13,32	7,19	0,17	4,87	2,98
ARPANET	1.05	2.83	1,84	1,09	5,40	2,57	4,42	8,49	5,99
BHVAC	0.24	1,16	0,65	0,09	4,35	0,94	3,93	8,30	5,78
EON	$\simeq 0$	3,37	2,31	$\simeq 0$	3,63	0,89	2,24	6,63	4,53
METRO	2,23	20,16	8.59	1,64	8,90	4,47	2,63	13,32	7,41
NSF	$\simeq 0$	12,97	3,29	0,04	4,43	1,76	0,15	7,46	2,21
PACBELL	1.89	5,27	3,61	0,32	3,00	1,44	2,30	8,00	4,76
VNSL	0,03	7,27	1.91	$\simeq 0$	8,57	3,07	0,09	4,92	1,34
<i>Global</i>	$\simeq 0$	20,16	3,26	$\simeq 0$	13,32	2,79	0,09	13,32	4,37

5.2.3 M/M/1 network cost function

The results obtained with this cost function are presented in Table 6. In this case, the best-response and the ACO algorithms perform better than GSA, for which relative errors up to 87.54% are obtained on some instances. The best-response algorithm achieves the lower average relative error with only 1%, whereas ACO has a lower maximum relative error with 14.99%. However, as can be noted from Table 7, the best-response algorithm provides the best tradeoff between optimization quality and computing times. Over the 800 instances, its worst execution time is only 4.7 seconds, whereas the computing times of GSA and ACO can be as high as 400 and 50 seconds, respectively.

Table 6: Relative gap (%) to the optimal multi-path solution for the M/M/1 cost function.

Topologie	BR			GSA			ACO		
	min	max	avg	min	max	avg	min	max	avg
ABOVENET	$\simeq 0$	1,52	0,72	$\simeq 0$	87,54	5,95	1,94	6,23	3,59
ARPANET	$\simeq 0$	0.50	0,13	$\simeq 0$	77,01	3,44	2,58	7,64	4,82
BHVAC	$\simeq 0$	1,45	0,35	$\simeq 0$	64,58	8,57	3,30	8,54	5,17
EON	$\simeq 0$	1.87	0,82	$\simeq 0$	58,09	3,15	1,71	8,92	3,63
METRO	$\simeq 0$	24.45	2.38	$\simeq 0$	30,89	3,62	0,20	14,99	2,77
NSF	$\simeq 0$	21.06	2.77	0,01	60,14	3,75	0,01	8,77	1,47
PACBELL	$\simeq 0$	1.19	0.20	$\simeq 0$	45,13	5,30	1,44	7,14	4,00
VNSL	$\simeq 0$	4.40	0.64	$\simeq 0$	24,07	3,51	0,34	6,82	2,46
<i>Global</i>	$\simeq 0$	24,45	1.00	$\simeq 0$	87,54	4,53	0,01	14,99	3,26

Table 7: Computing times (seconds) for the 800 problem instances.

Scenario	Quadratic Cost			M/M/1 Cost		
	BR	GSA	ACO	BR	GSA	ACO
Min	0,04	1,02	4,73	0,03	0,75	4,98
Max	6.75	297,73	48,83	4.68	400,82	50,95
Average	0,92	30,85	22,16	0,72	27,45	22,40

Remark 2. *The results obtained with the standard best-response algorithm are comparable to those presented above for the penalized best-response algorithm. In practice, the results of the standard best-response algorithm are sometimes better, and sometimes worse, but the difference is always in the order of a few percents. On the other hand, the computing times of the standard best-response algorithm are clearly better since its worst execution time over the 800 instances is 1.1 seconds for the M/M/1 cost. This is something expected since the penalty term implies an extra computation with respect to the standard best-response algorithm.*

6 Conclusion

In this paper, we studied non-linear single path routing problems, which are known to be extremely hard to solve. Yet, these problems naturally arise when one seeks to optimize a quality of service metric such as the delay. For solving these problems, we have proposed an algorithm inspired from Game Theory in which individual flows are allowed to independently select their path to minimize their own cost function. Given the routing

of the other flows, each flow seeks to minimize the sum of its own end-to-end delay, plus a penalty term which plays the role of an incentive for players to take into account the impact of their actions on the others. We have proven the convergence of this algorithm to a Nash equilibrium of the game. In the case of polynomial link latency functions, we have also established a tight upper bound on the approximation ratio of the algorithm. The numerical results obtained with this algorithm are quite interesting: the relative gap to the optimal multi-path solution remains modest (equivalent to that of GSA and ACO), while the execution times are substantially lower than that of the other algorithms. These results suggest that the best-response algorithm is an appropriate solution for large-scale single-path routing problems with nonlinear cost functions, where other techniques show their limits.

Acknowledgement

This work was supported by French FUI project NEC (AAP FUI 12).

References

- [1] A Al Sheikh, O. Brun, P.E. Hladik, and B.J. Prabhu. A best-response algorithm for multiprocessor periodic scheduling. In *Real-Time Systems (ECRTS), 2011 23rd Euromicro Conference on*, pages 228–237, July 2011.
- [2] F. Alvelos and J.M. Valerio de Carvalho. Comparing branch-and-price algorithms for the unsplittable multicommodity flow problem. In *Proceedings of INOC'2003*, pages 7–12, Evry-Paris, 2003.
- [3] Baruch Awerbuch, Yossi Azar, and Amir Epstein. The price of routing unsplittable flow, 2005.
- [4] C. Barnhart, C. Hane, and P. Vance. Using branch-and-price-and-cut to solve origin-destination integer multicommodity flow problems. *Operations Research*, 48(2):318–326, 2000.
- [5] M. Belaidouni and W. Ben-Ameur. A superadditive approach to solve the minimum cost single path routing problem: preliminary results. In *Proceedings of INOC'2003*, pages 67–71, Evry-Paris, France, 2003.
- [6] M. Belaidouni and W. Ben-Ameur. On the minimum cost multiple-source unsplittable flow problem. *RAIRO-Operations Research*, 2005.
- [7] Kshipra Bhawalkar, Martin Gairing, and Tim Roughgarden. Weighted congestion games: The price of anarchy, universal worst-case examples, and tightness. *ACM Trans. Economics and Comput.*, 2(4):14:1–14:23, 2014.

- [8] Erik Gunnar Boman. *Infeasibility and Negative Curvature in Optimization*. PhD thesis, Stanford, CA, USA, 1999. AAI9924533.
- [9] E. Cela. *The Quadratic Assignment Problem: Theory and Algorithms*. Kluwer Academic Publishers, 1998.
- [10] Pierre Coucheney. *Auto-optimisation des réseaux sans fil : une approche par la théorie des jeux*. PhD thesis, Université de Grenoble, 2006.
- [11] B. Davie and Y. Rekhter. *MPLS Technology and Applications*. Morgan Kaufmann, 2000.
- [12] Marco Dorigo. Optimization, learning and natural algorithms. *Ph. D. Thesis, Politecnico di Milano, Italy*, 1992.
- [13] Marco Dorigo, Vittorio Maniezzo, and Alberto Colorni. The ant system: Optimization by a colony of cooperating agents. *26th IEEE Transactions on Systems, Man, and Cybernetics - part B*, (1):29–41, 1996.
- [14] M. Junger et al., editor. *50 years of Integer Programming 1958–2008*. Springer Berlin, 2010.
- [15] B. Fortz and M. Thorup. Internet traffic engineering by optimizing ospf weights. In *Proceedings of IEEE InfoCom 2000, 19th Annual Joint Conference of the IEEE Computer and Communications Societies*, volume 2, pages 519–528, Tel-Aviv, Israel, March 2000. IEEE.
- [16] D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish unsplittable flows. *Theoretical Computer Science*, (348):226–239, 2005.
- [17] Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, Cambridge, MA, 1991.
- [18] J.M. Garcia, A. Rachdi, and O. Brun. Optimal lsp placement with qos constraints in diffserv/mpls networks. In *Proceedings of the 18th International Teletraffic Congress - ITC-18*, volume 5 of *Teletraffic Science and Engineering*, pages 11 – 20. 2003.
- [19] J. Geffard. A solving method for singly routed traffic demands in telecommunication networks. *Annales des Télécommunications*, 56(3-4):140–149, 2001.
- [20] Leonard Kleinrock. *Queueing Systems*, volume I: Theory. Wiley Interscience, 1975.
- [21] Walter Murray and Kien-Ming Ng. An algorithm for nonlinear optimization problems with binary variables. *Computational Optimization and Applications*, 47(2):257–288, 2010.

- [22] Panos M. Pardalos. Continuous approaches to discrete optimization problems. In G. Di Pillo and F. Giannessi, editors, *Nonlinear Optimization and Applications*, pages 313–325. Springer US, 1996.
- [23] S. Park, D. Kim, and K. Lee. An integer programming approach to the path selection problems. In *Proceedings of INOC'2003*, pages 448–453, Evry-Paris, 2003.
- [24] M. Pióro and D. Medhi. *Routing, flow, and capacity design in communication and computer networks*. Elsevier, 2004.
- [25] R. W. Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal on Game Theory*, 2:65–67, 1973.
- [26] Ruud Schoonderwoerd, Owen Holl, Janet Bruten, and Leon Rothkrantz. Ant-based load balancing in telecommunications networks. *Adaptive Behavior*, 5:169–207, 1996.
- [27] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson. Measuring isp topologies with rocketfuel. *Networking, IEEE/ACM Transactions on*, 12(1):2 – 16, feb. 2004.
- [28] M. Wang, C. W. Tan, A. Tang, and S. H. Low. How bad is single-path routing? In IEEE, editor, *GlobeCom 2009*, 2009.

A Proof of theorem 3

Before proving theorem 3, let us first introduce some notations. Let $\boldsymbol{\pi}$ be a solution computed with the best-response algorithm and $\boldsymbol{\pi}^*$ be an optimal solution. In the following, to simplify notations, we denote by y_e and y_e^* the traffic volumes $y_e(\boldsymbol{\pi})$ and $y_e(\boldsymbol{\pi}^*)$, respectively. Since $\boldsymbol{\pi}$ is a Nash equilibrium, we have $c_i(\boldsymbol{\pi}) \leq c_i(\boldsymbol{\pi}_i^*, \boldsymbol{\pi}_{-i})$, for all $i \in \mathcal{K}$, which according to lemma 1 translates into $F(\boldsymbol{\pi}) \leq F(\boldsymbol{\pi}_i^*, \boldsymbol{\pi}_{-i})$ for all i . The deviation of player i from path π_i to path π_i^* partitions the set of links in three separate subsets: those for which the load increases ($e \in \pi_i^* \setminus \pi_i$), those for which it decreases ($e \in \pi_i \setminus \pi_i^*$), and those for which it remains constant. Let us therefore define

$$\alpha_e^i = \begin{cases} 1 & \text{if } e \in \pi_i^* \setminus \pi_i, \\ 0 & \text{otherwise.} \end{cases} \quad \text{and} \quad \beta_e^i = \begin{cases} 1 & \text{if } e \in \pi_i \setminus \pi_i^*, \\ 0 & \text{otherwise.} \end{cases}$$

In order to simplify the presentation of the proof, we shall first prove the result for monomials, and then generalize to arbitrary polynomials. The monomial case contains most of the essential ingredients of the proof and has the added advantage of being notationally lighter. Note that in the course of the proof, we shall use two technical results which are proved at the end of this appendix.

Proposition 6. *If $\ell_e(x) = a_e x^d$, then the approximation factor of the penalized best-response algorithm is upper bounded by $(2^{1/(d+1)} - 1)^{-(d+1)}$.*

Proof. For all $i \in \mathcal{K}$, the condition $F(\boldsymbol{\pi}) \leq F(\boldsymbol{\pi}_i^*, \boldsymbol{\pi}_{-i})$ can be written as follows

$$0 \leq \sum_e a_e \left((y_e + (\alpha_e^i - \beta_e^i) \lambda_i)^{d+1} - y_e^{d+1} \right) \quad (24)$$

$$= \sum_e a_e \left(\alpha_e^i \lambda_i \sum_{k=0}^d (y_e + \alpha_e^i \lambda_i)^k y_e^{d-k} - \beta_e^i \lambda_i \sum_{k=0}^d (y_e - \beta_e^i \lambda_i)^k y_e^{d-k} \right) \quad (25)$$

$$= \sum_e a_e \left(\alpha_e^i \lambda_i \sum_{k=1}^d (y_e + \alpha_e^i \lambda_i)^k y_e^{d-k} - \beta_e^i \lambda_i \sum_{k=1}^d (y_e - \beta_e^i \lambda_i)^k y_e^{d-k} \right) + \sum_e a_e (\alpha_e^i - \beta_e^i) \lambda_i y_e^d \quad (26)$$

$$\leq \sum_e a_e \left(\alpha_e^i \lambda_i \sum_{k=1}^d (y_e + \alpha_e^i \lambda_i)^k y_e^{d-k} \right) + \sum_e a_e (\alpha_e^i - \beta_e^i) \lambda_i y_e^d \quad (27)$$

$$\leq \sum_e a_e \left(\alpha_e^i \lambda_i \sum_{k=1}^d (y_e + y_e^*)^k y_e^{d-k} \right) + \sum_e a_e (\alpha_e^i - \beta_e^i) \lambda_i y_e^d, \quad (28)$$

where equality (25) follows from lemma 2 with $\gamma = \alpha_e^i \lambda_i$ and $\delta = \beta_e^i \lambda_i$. To get to (26), we have separated the $k = 0$ term in each of the two sums and combined them together.

Summing the inequalities (28) over i and using the fact that $\sum_i (\alpha_e^i - \beta_e^i) \lambda_i = y_e^* - y_e$ yields

$$0 \leq \sum_e a_e y_e^* \sum_{k=1}^d (y_e + y_e^*)^k y_e^{d-k} + \sum_e a_e y_e^* y_e^d - \sum_e a_e y_e^{d+1} \quad (29)$$

$$= \sum_e a_e y_e^* \sum_{k=0}^d (y_e + y_e^*)^k y_e^{d-k} - \sum_e a_e y_e^{d+1} \quad (30)$$

$$= \sum_e a_e (y_e + y_e^*)^{d+1} - \sum_e a_e y_e^{d+1} - \sum_e a_e y_e^{d+1} \quad (31)$$

$$= \sum_e a_e (y_e + y_e^*)^{d+1} - 2 \sum_e a_e y_e^{d+1} \quad (32)$$

$$\leq (s^{1/(d+1)} + (s^*)^{1/(d+1)})^{d+1} - 2 \sum_e a_e y_e^{d+1}, \quad (33)$$

where the last inequality follows from lemma 3 with $s = \sum_e a_e y_e^{d+1}$ and $s^* = \sum_e a_e (y_e^*)^{d+1}$. Rearranging the above inequality, we get

$$2F(\boldsymbol{\pi}) \leq \left(F(\boldsymbol{\pi})^{1/(d+1)} + F(\boldsymbol{\pi}^*)^{1/(d+1)} \right)^{d+1}, \quad (34)$$

from which we deduce that

$$\frac{F(\boldsymbol{\pi})}{F(\boldsymbol{\pi}^*)} = \left(\frac{1}{2^{1/(d+1)} - 1} \right)^{d+1}, \quad (35)$$

as was stated in the proposition. \square

Proof of theorem 3. For each monomial of degree j , we shall follow the same steps as for the monomial case until inequality (33) to arrive at

$$0 \leq \sum_j \left(s_j^{1/(j+1)} + (s_j^*)^{1/(j+1)} \right)^{j+1} - 2 \sum_j \sum_e a_{e,j} y_e^{j+1}, \quad (36)$$

where $s_j = \sum_e a_{e,j} y_e^{j+1}$ and $s_j^* = \sum_e a_{e,j} (y_e^*)^{j+1}$. Since $(x_1^p + x_2^p)^{1/p} \leq (x_1^q + x_2^q)^{1/q}$ for $p \geq q > 0$, the above inequality becomes

$$2F(\boldsymbol{\pi}) \leq \sum_j \left(s_j^{1/(d+1)} + (s_j^*)^{1/(d+1)} \right)^{d+1}. \quad (37)$$

Applying lemma 3 to the sum on the right-hand side, we get

$$2F(\boldsymbol{\pi}) \leq \left(F(\boldsymbol{\pi})^{1/(d+1)} + F(\boldsymbol{\pi}^*)^{1/(d+1)} \right)^{d+1}, \quad (38)$$

from which the stated result follows. \square

Lemma 2 (Decomposition). *Let γ and δ be two non-negative numbers such that $\gamma\delta = 0$. Let m be a non-negative integer and $x \geq 0$. Then,*

$$(x + (\gamma - \delta))^{m+1} - x^{m+1} = \gamma \sum_{k=0}^m (x + \gamma)^k x^{m-k} - \delta \sum_{k=0}^m (x - \delta)^k x^{m-k}. \quad (39)$$

Proof. The proof is based on simple algebraic manipulations but we give it for the sake of completeness. Starting from the left-hand side, we have

$$(x + (\gamma - \delta))^{m+1} - x^{m+1} = (\gamma - \delta) \sum_{k=0}^m (x + (\gamma - \delta))^m x^{m-k} \quad (40)$$

$$= \gamma \sum_{k=0}^m \left(\sum_j \binom{k}{j} (x + \gamma)^j (-\delta)^{k-j} \right) x^{m-k} \quad (41)$$

$$- \delta \sum_{k=0}^m \left(\sum_j \binom{k}{j} (x - \delta)^j \gamma^{k-j} \right) x^{m-k} \quad (42)$$

$$= \gamma \sum_{k=0}^m (x + \gamma)^k x^{m-k} - \delta \sum_{k=0}^m (x - \delta)^k x^{m-k}, \quad (43)$$

where the last equality follows from the assumption that $\gamma\delta = 0$. \square

Lemma 3. For non-negative numbers a_e , x_e , z_e , and non-negative integer m ,

$$\sum_e a_e (x_e + z_e)^{m+1} \leq \left(s_x^{1/(m+1)} + s_z^{1/(m+1)} \right)^{m+1} \quad (44)$$

where $s_x = \sum_e a_e x_e^{m+1}$ and $s_z = \sum_e a_e z_e^{m+1}$.

Proof. The proof is an application of Hölder's inequality. Starting from the left-hand side,

$$\sum_e a_e (x_e + z_e)^{m+1} = \sum_e a_e \sum_{k=0}^m \binom{m+1}{k} x_e^k z_e^{m+1-k} \quad (45)$$

$$= \sum_{k=0}^{m+1} \binom{m+1}{k} \left(\sum_e a_e x_e^k z_e^{m+1-k} \right) \quad (46)$$

$$= \sum_{k=0}^{m+1} \binom{m+1}{k} \left(\sum_e (a_e x_e^{m+1})^{k/(m+1)} (a_e z_e^{m+1})^{(m+1-k)/(m+1)} \right) \quad (47)$$

$$\leq \sum_{k=0}^{m+1} \binom{m+1}{k} \left(\sum_e a_e x_e^{m+1} \right)^{k/(m+1)} \left(\sum_e a_e z_e^{m+1} \right)^{(m+1-k)/(m+1)} \quad (48)$$

$$= \sum_{k=0}^m \binom{m+1}{k} s_x^{k/(m+1)} s_z^{(m+1-k)/(m+1)} \quad (49)$$

$$= \left(s_x^{1/(m+1)} + s_z^{1/(m+1)} \right)^{m+1}, \quad (50)$$

where we have used the Hölder's inequality with $p = (m+1)/k$ and $q = (m+1)/(m+1-k)$ to get to (48) from (47). \square