

A Framework for Improving Routing Configurations using Multi-Objective Optimization Mechanisms

Pedro Sousa, Vítor Pereira, Paulo Cortez, Miguel Rio, and Miguel Rocha

Abstract—IP networks are nowadays well established technologies being used to support a myriad of applications and services, thus assuming a crucial role in today's telecommunication systems. Nevertheless, such infrastructures usually require network administrators to perform a wide set of complex planning and management tasks trying to attain adequate network configurations. Many of such management tasks can be mathematically formulated as NP-hard optimization problems, sometimes involving several objective functions. In this context, this work explores and demonstrates the potential of using computational intelligence methods as optimization engines to tackle complex network optimization problems. In particular, Multi-objective Evolutionary Algorithms (MOEAs) are used to attain near-optimal link state routing configurations robust to distinct operational conditions. As result, network administrators will be provided with a set of alternative routing configurations representing distinct tradeoffs between the considered optimization goals. The robustness of the proposed methods is illustrated by presenting several multi-objective optimization examples able to improve the performance and resilience levels of a network infrastructure.

In this perspective, this work presents a contribution for this research area by proposing specific MOEAs based optimization methods able to improve network routing configurations. Furthermore, the devised methods are also integrated in a freely available Traffic Engineering optimization framework able to be used by network administrators interested in improving the routing configurations of their network infrastructures.

Index Terms—Communications Software, Routing, Traffic Engineering, Network Resilience, Multi-Objective Evolutionary Algorithms

I. INTRODUCTION

Nowadays, IP based network infrastructures have to support a myriad of applications and services generating high volumes of traffic, and many of them with strict operational and availability requirements. In this perspective, actual network infrastructures should present high levels of resilience in order to behave correctly and efficiently under a wide set of operational conditions [1]. As it is well known, routing protocols are key elements of IP converged networks, thus

having a major influence in the operational conditions of such communication infrastructures. In this specific field, link-state routing approaches such as Intermediate System to Intermediate System (IS-IS) [3] or Open Shortest Path First (OSPF) [2] protocols are very popular, being often used by Internet Service Providers (ISPs) administrators to deliver connectivity between all network equipment. In these protocols, each network link is assigned with a specific cost value from which the shortest-paths are computed.

In this perspective, an optimized routing configuration is essential to assure that an efficient traffic distribution is attained in the network infrastructure. Otherwise, and depending on the required traffic demands, it is possible that specific parts of the network might be congested, or at least with very high utilization rates, while some topology links could be almost unused. Furthermore, on such circumstances, specific link failures that may occur in the network may have a huge impact on the degradation of the infrastructure service levels [4][5]. Such inefficient traffic distribution in the network infrastructure has several important consequences. Firstly, an inefficient use of network resources represents a cost penalty from the ISP economical perspective. Secondly, the existence of congested links in the network affects the performance of the communication infrastructure, decreasing the offered service quality and imposing possible violations to the Service Level Agreements (SLA) negotiated with specific ISP customers. Finally, a well established and adequate routing infrastructure is essential to underpin the efficient deployment of complementary finer-grain Quality of Service (QoS) oriented mechanisms [8] in the network (e.g. traffic differentiation mechanisms, resource reservation approaches, admission control schemes, etc.).

The Traffic Engineering (TE) area addressed by this work is specially focused on the performance evaluation and optimization of operational IP networks. In this specific context, some relevant research focused on the specific objective of achieving an efficient traffic distribution in the networking infrastructures, taking into account the expected traffic demands. Such traffic demands are usually expressed by traffic matrices (e.g. [9], [21]) which are used by ISPs to better understand the traffic aggregates traversing the network domain. Thus, by combining the traffic demands and the traffic forwarding dynamics assumed by the routing protocols, it is possible to foresee which links of the network topology might be overloaded.

As an example, some studies conducted in this area focus on the configuration optimization of commonly used routing

Manuscript received June 14, 2016; revised September 30, 2016.

Pedro Sousa and Vítor Pereira are with the Centro ALGORITMI and Department of Informatics, University of Minho, Braga, Portugal (Emails: {pns@di, Vitor.Pereira@algoritmi}.uminho.pt).

Paulo Cortez is with the Centro ALGORITMI and Department of Information Systems, University of Minho, Guimarães, Portugal (Email: pcortez@dsi.uminho.pt).

Miguel Rio is with the Department of Electric and Electronic Engineering, University College London, London, U.K. (Email: m.rio@ee.ucl.ac.uk).

Miguel Rocha is with the Centre of Biological Engineering and Department of Informatics, University of Minho, Braga, Portugal (email: mrocha@di.uminho.pt).

protocols (e.g. [6], [7]) while others propose new routing approaches to attain more efficient traffic distribution (e.g. [22]) in ISP networks. Many of such routing optimization approaches usually translate to NP-hard optimization problems that seek to find a set of routing weights that are able to optimize the congestion levels of the network, considering specific aggregated traffic demands. For this specific purpose, the use of computational intelligence methods to solve TE related problems has presented encouraging results and Evolutionary Algorithms (EAs) have been successfully used to solve congestion based formulations, or other variants involving multi-constrained optimization approaches ([14], [10], [11], [12], [13]). In addition, other meta-heuristics were also used in such TE weight setting optimization problems, namely techniques such as Local Search and Simulated Annealing [15], [16]. Most of such routing optimization proposals assume Equal-cost multi-path routing (ECMP) behavior to distribute traffic among several equal-cost routes to a given destination. Nevertheless, other approaches also focus on the use of a single shortest path between a given network source and destination [17] or, based in the work present in [18], also tackle such TE problem assuming that non-shortest paths may also be used to improve load-balancing levels and achieve a better distribution of traffic within the network infrastructure [19]. In a global perspective, in [20] the authors presents an extensive study overviewing several TE optimization techniques and approaches that have been deployed for managing intra-domain routing in networks operated with shortest path routing protocols.

Considering the above mentioned, this work foster this research field by presenting a contribution specifically focused on devising mechanisms able to provide enhanced routing configurations, using Multi-objective Evolutionary Algorithms (MOEAs) as the main optimization engines. In particular, the proposed formulations are able to improve routing configurations in order to support multi constrained QoS scenarios and also attaining resilient network infrastructures. As a complementary contribution this work presents an automated and intelligent optimization framework integrating the mentioned objectives. The framework is able to assist network administrators in the configuration of resilient network infrastructures, providing them with a set of configuration alternatives expressing distinct trade-offs between the considered objectives. To illustrate the capabilities of the devised framework, three optimization scenarios are presented illustrating the TE framework capabilities in providing enhanced routing configurations.

The remainder of the document is organized as follows. Section II presents the traffic engineering framework developed in the context of this work and its main components. Next, section III focuses on the mathematical formulation adopted by the framework, highlighting the role of Evolutionary Computation methods in the optimization approaches. Section IV presents some illustrative optimization scenarios dealing with network QoS and resilience issues, also depicting the obtained results of the proposed methods. Finally, section V presents the conclusions of the work and comments on the computational efforts and availability issues of the devised framework.

II. A TRAFFIC ENGINEERING OPTIMIZATION FRAMEWORK

This section briefly describes the Traffic Engineering (TE) framework developed by the authors which resorts to MOEAs as optimization engines. This freely available framework allows network administrators to perform several TE optimization operations, thus being able to optimize the behavior of their own network infrastructures. Furthermore, it could also be used to devise, implement and test new TE approaches taking advantage of evolutionary computation inspired optimization engines. Figure 1 depicts a conceptual view of the developed framework which internal modules are now summarized.

A. Framework Modules

OSPF routing simulator: This module implements an OSPF routing simulator taking as inputs the network topology and a specific routing configuration (i.e. a set of OSPF link costs). The link costs can be manually assigned by the administrator, automatically assigned by a set of available heuristics, or as a result of a specific optimization process. Based on the link costs it is possible to calculate the edge-to-edge paths that will be used in the network (e.g. using the Dijkstra algorithm [23]). Using the calculated paths and the considered traffic demands, this module will then distribute the traffic in the network infrastructure, being possible to assess the quality of the considered routing configuration using appropriate cost functions. The mathematical model sustaining part of the operations described in this module is detailed in section III-A.

Computational intelligence optimization: This framework core module aims to attain near-optimal routing configuration for the considered scenario. This module integrates several multi-objective optimization mechanisms, e.g. Non-dominated Sorting Genetic Algorithm (NSGA-II) and Strength Pareto Evolutionary Algorithm (SPEA2), that are later described (section III-B). This module also receives the fitness functions that will be used to measure the routing solutions quality for a given optimization scenario. Examples of such functions will be given in section IV when describing illustrative optimization scenarios.

Topology, traffic demands and failures inputs: Several inputs are assumed by the framework depending on the considered scenarios. In real use case scenarios, it is expected that network administrators provide the framework with a representation of the network topology and corresponding characteristics (e.g. a graph based representation) to apply the optimization engines supported by the framework.

In this work, to test and assess the presented multi-objective optimization formulations, a set of synthetic networks is used. The instances were generated by using the Brite topology generator [27], varying the number of nodes ($N=30, 50$) and the average degree of each node ($m=2, 3, 4$), which resulted in a set of network instances trying to model real network infrastructures. The link bandwidth (capacity) varies between 1 and 10 Gbits/s under a uniform distribution. The networks were generated using the Barabasi-Albert model, using a heavy-tail

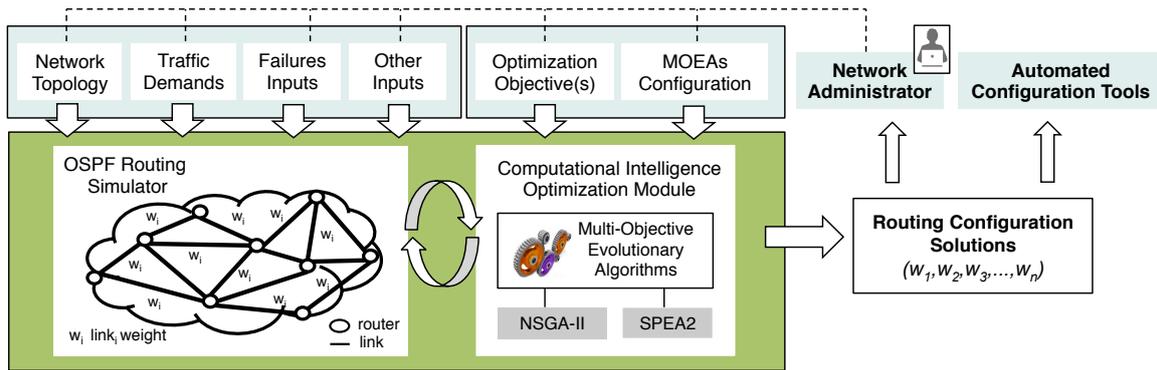


Fig. 1. Modules integrating the Traffic Engineering framework assisted by Computational Intelligence methods.

distribution and an incremental grow type (parameters HS and LS were set to 1000 and 100, respectively).

Traffic demand matrices are also inputs to the framework. Such matrices express the global amounts of traffic that is expected to traverse the network infrastructure. ISPs often estimate traffic demand matrices to improve resource management tasks and there are several techniques to obtain such matrices ([9], [21]). A traffic demand matrix usually summarizes, for each source/destination edge router pair, a given bandwidth required to be supported by the network domain. In this context, when using synthetic networks, the framework is also able to tune the difficulty levels of the optimization problem, by considering distinct levels of traffic demands.

Optimization objectives and MOEAs configuration:

These inputs define the optimization objectives that guide the intelligent optimization methods (some examples are given in section IV) and several fine tune configuration parameters for the optimization mechanisms.

Routing solutions: As observed in Figure 1, the output of the framework is a set routing configurations representing distinct tradeoffs between the considered objectives. Such set of routing solutions is presented to the network administrator as a Pareto front (details about this topic are explored in section III-B).

B. Framework User Interfaces

The focus of this paper is on the description of the internal optimization mechanisms and problem formulations assumed by the framework to provide enhanced routing configurations. However, it is important to highlight that the devised optimization framework provides an easy and intuitive interface to network administrators.

As example, Figure 2 shows some user interfaces of the framework, namely a representation of the network topology which is being optimized, a set of input parameters provided by the user to guide the optimization process and a set of computed routing configurations with the corresponding performance metrics. Thus, independently of the complexity of the mechanisms and problem formulations used by the framework and which are described in the remain of the paper, users can use the framework in an easy and intuitive way.

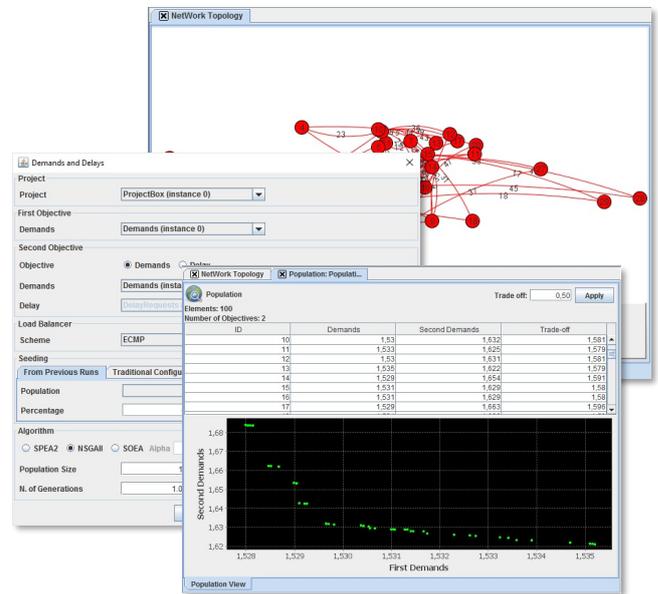


Fig. 2. Examples of user interfaces of the devised optimization framework.

III. MATHEMATICAL FORMULATION AND MULTI-OBJECTIVE EVOLUTIONARY ALGORITHMS

This section presents the mathematical model assumed by the framework to represent the network infrastructure, some details regarding the use of the MOEAs and comparative heuristics for routing configuration.

A. Mathematical Model

We start by defining the generic mathematical model assumed by the framework. Based on such model, some variants are later introduced for each one of the illustrative scenarios.

Most of the selected network optimization methods aim to attain efficient distribution of the traffic aggregates in the links of the network domain. Thus, as much as possible, the model aims to avoid the existence of congested links in the infrastructure. The framework represents the network topology as a direct graph $G(N, A)$, with N representing a set of nodes (network routers), and A representing a set of arcs (network links), with a capacity of c_a for each $a \in A$.

For a specific routing configuration, and considering a specific traffic matrix given as input (D), $f_a^{(s,t)}$ expresses the amount of traffic routed over the arc a having source s and destination t . Thus, the utilization of an arc a can be defined as in Eq. 1 with ℓ_a being the sum of all flows $f_a^{(s,t)}$ that travel over it. Considering the utilization degree of an arc, we adopt the cost function, Φ_a whose derivative is presented by Eq. 2, as a linear cost function which penalizes high congested links (a well known function in the TE area proposed by Fortz and Thorup [26]).

$$u_a = \frac{\ell_a}{c_a} \quad (1)$$

$$\Phi'_a = \begin{cases} 1 & \text{for } 0 \leq u_a < \frac{1}{3} \\ 3 & \text{for } \frac{1}{3} \leq u_a < \frac{2}{3} \\ 10 & \text{for } \frac{2}{3} \leq u_a < \frac{9}{10} \\ 70 & \text{for } \frac{9}{10} \leq u_a < 1 \\ 500 & \text{for } 1 \leq u_a < \frac{11}{10} \\ 5000 & \text{for } u_a \geq \frac{11}{10} \end{cases} \quad (2)$$

$$\Phi = \sum_{a \in A} \Phi_a \quad (3)$$

Considering the above mentioned mathematical representation, a possible optimization objective is to try to distribute traffic demands in the network to minimize the sum of all costs, as expressed by Equation 3. To be possible to compare the results obtained in distinct topologies, the congestion measure can be normalized over distinct topology scenarios, by using a scaling factor which considers the minimum hop count between every pair of nodes and the individual demands expressed in the demand matrix (as mentioned in [13]). Thus, in this work, we use a normalized congestion measure, entitled as Φ^* , as the main function used to assess the quality of a given routing network configuration. A detailed analysis of the Φ^* function allows to conclude that when Φ^* equals 1, all loads are below $1/3$ of the link capacity, while when all arcs are exactly full the value of Φ^* is $10/3$. In the results presentation, this value is considered as a threshold that bounds the acceptable working region of the network.

As mentioned, the illustrative scenarios presented in this work deal with multi-objective optimization problems, which means that they target the simultaneous minimization of several Φ^* functions under distinct operational conditions, or combinations involving the function Φ^* and other defined functions of distinct nature. More details regarding such approaches are given in Section IV along with the details related to the generation of distinct traffic demand matrices (D) to tune the difficulty level of the optimization problems.

B. Multi-objective Evolutionary Algorithms

From the mid-1980s that MOEAs are used to solve several multiple-optimization problems, being one of the most competitive approaches in this field [28]. The devised optimization framework integrates two popular algorithms in this field, the SPEA2 [24] and the NSGA-II [25], which are two of the algorithms with best performance.

The developed optimization framework adapts the MOEA optimization engines to the particular case of optimizing routing configurations. Thus, in the used MOEAs, each individual encodes a routing solution, being a vector of integer values, where each value (gene) corresponds to the weight of a link (arc) in the network. Thus, the size of the individual equals the number of links in the network infrastructure. In real implementations, OSPF link weights are integers valued from 1 to 65535, but here only values in range [1; 20] were considered, allowing to reduce the search space and, simultaneously, increasing the probability of finding equal cost multipaths, which greatly benefits effective bandwidth use by allowing load-balancing traffic schemes over multiple paths.

In the MOEA optimization process, the initial population is filled with randomly generated individuals, where the arc weights are taken from a uniform distribution within the allowed range. To generate new individuals, the MOEAs use several reproduction operators allowing the individuals recombination and maintaining genetic diversity:

- Random mutation: replaces a given gene by a random value, within the allowed range;
- Incremental/decremental mutation: replaces a given gene by the next or by the previous integer value, with equal probabilities, within the allowed range;
- Uniform crossover: this operator works by taking two parents as input and generating two offspring. For each position in the genome, a binary variable is randomly generated: if its value is 1, the first offspring takes the gene from the first parent in that position, while the second offspring takes the gene from the second parent; if the random value is 0, the roles of the parents are reversed.

The quality of each individual (i.e. of a specific routing configuration) is then evaluated in accordance with the considered objective functions (e.g. Φ^* functions or other depending on the formulated problem) and the best individuals proceed to the next optimization rounds. After a given number of iterations the individuals with the best performance are selected to be presented to the user.

The MOEAs mechanisms are appropriate to deal with the multi-objective nature of the optimization problems discussed in the following sections. In fact, MOEAs return a set of solutions with distinct trade-offs between the considered objectives, allowing the network administrator to decide which solution to implement. Figure 3 presents a conceptual view of the solutions (Pareto front) returned by the optimization framework assuming an optimization process trying to minimize two objective functions. In this case, the quality of a specific solution is measured using *Function X* and *Function Y*, obtaining fitness values with *x-value* and *y-value*, respectively for each one of the objectives. The considered functions will measure the performance of the network when considering different objectives, as will be later illustrated. Each of the returned solutions assumes a given tradeoff between the considered objectives, which makes easier the selection of the most appropriate one for a given scenario. As illustrated in Figure 3, the network administrator (or some type of automated

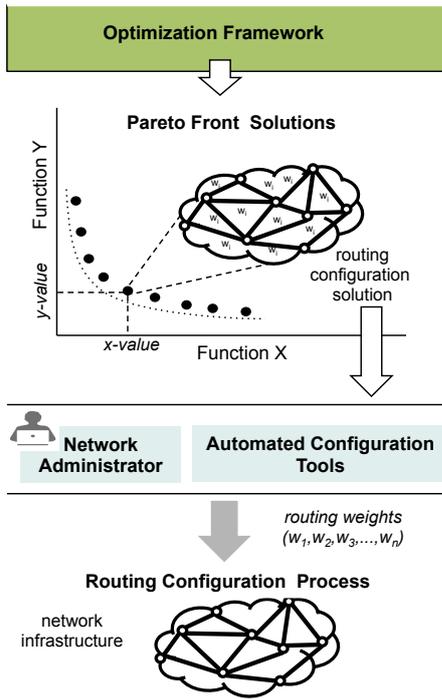


Fig. 3. Conceptual example of the framework optimization outputs.

configuration tool) will select the most appropriate routing solution and proceed to the configuration of the network infrastructure. Thus, in Figure 3, each point from the Pareto front is in fact a set of routing weights (w_1, w_2, \dots, w_n) that can be assigned to the network links of the topology.

C. Comparative Scenarios

To assess the quality of the obtained results from the multi-objective approach, a number of traditional heuristic methods were also implemented and included in the framework. In the given optimization scenarios, the results of some of these heuristics were also presented for a better understanding of the advantages of the proposed mechanisms. The included heuristics are:

- **InvCap:** this heuristic sets each link weight to a value inversely proportional to its capacity. This method is usually assumed in the OSPF protocol Cisco implementations, as a default mechanism to assign OSPF costs to each link. With this heuristic, links with higher capacities are more prone to be included in the computed shortest paths, being a simple way to try to improve the network performance.
- **L2:** sets each link weight to a value proportional to its Euclidean distance. This metric can be useful for comparative purposes, when some of the optimization objectives also involve the propagation delays of the paths. With this heuristic, lower weights are given to links with lower propagation delays, being more prone to be included in the computed shortest paths. This may be useful for comparison purposes when the minimization of network delays integrates the optimization objectives.

TABLE I
SOME ILLUSTRATIVE OPTIMIZATION SCENARIOS

Scenario	General Description
Multiconstrained QoS Optimization [Section IV-A]	The framework is used to simultaneously optimize two objectives related with QoS (Quality of Service) metrics. One of the objectives is to attain an efficient distribution of the traffic in the network links, thus avoiding possible link congestion events. The other envisaged objective is to assure that specific delay constraints are accomplished in specific edge-to-edge paths of the network infrastructure.
Resilience to Traffic Variations [Section IV-B]	The framework is used to attain routing configurations able to assure resilience levels to variations in the traffic volumes traversing the network infrastructure. The suggested configurations should be able to support a heterogeneous set of traffic volumes that may vary during a given time period.
Resilience to Link Failures [Section IV-C]	The administrator is interested in improving the resilience levels when dealing with failures of specific links of the network infrastructure. Thus, the objective is that even in the presence of such failure events, the network infrastructure still presents acceptable performance with minimal traffic loss.

- **Random:** a number of randomly generated solutions (i.e. link weights) are analyzed and the best solution is selected. The number the random solutions analyzed is usually equal to the number of tested solutions by the evolutionary algorithms. With this comparative heuristic the objective is to prove the usefulness of the used intelligent optimization mechanisms, by comparing the obtained solution with the best from the randomly generated ones.
- **Unit:** sets every link weight to one (i.e. best paths are the ones with the minimum number of hops). This is a common strategy for routing configurations, when the objective is to establish paths involving the minimum number of nodes between source/destination nodes.

IV. FRAMEWORK ILLUSTRATIVE OPTIMIZATION RESULTS

In this section, we present several illustrative examples of the framework optimization capabilities. Table I summarizes the three selected illustrative scenarios.

In the following sections, the results were obtained with the framework optimization module operating with the NSGA-II algorithm.

A. Multiconstrained QoS Optimization

This optimization scenario focus on a multi-constrained optimization of routing weights involving two distinct objectives. One objective is to attain a balanced and efficient distribution of traffic into the network avoiding as far as possible the existence of congested links. The second objective is to obey

to specific propagation delays limits between specific pairs of edge-to-edge routers in the topology. Such constraints express that the ISP wants to force lower delays between specific node pairs, which might be motivated by diverse reasons (e.g. such pair of nodes may represent access points of a high priority customer of the ISP; the paths connecting such node pairs are expected to be traversed by delay sensitive traffic, etc.)

Therefore, the framework assumes that client demands are mapped into a matrix summarizing, for each source/destination edge router pair, a given required bandwidth and, if also defined, a target edge-to-edge delay to be supported by the network domain.

1) *Problem Formulation*: In this problem formulation (see Table II), given a specific network topology, a traffic demand matrix (D) and a delay requirements matrix (DR), the objective is to achieve a set of OSPF weights (w) that simultaneously minimize the functions $\Phi^*(w)$ and $\gamma^*(w)$, which are the penalty functions for congestion and edge-to-edge delays, respectively. The $\gamma^*(w)$ function is evaluated in a similar way to the $\Phi^*(w)$ function (section III-A), but now considering the propagation delays targets specified for the paths established between specific pairs of routers (expressed by the DR matrix) and the delay values effectively obtained for a particular routing weight configuration.

Thus, the cost associated with a given solution w is evaluated using functions $\Phi^*(w)$ for congestion and $\gamma^*(w)$ for delays, with both functions normalized in the same range. As explained before, acceptable values for such network functions are in the range $[1, 10]$ meaning that the traffic demands and delay restrictions (D , DR matrices) are accomplished by the routing configuration.

In order to generate the traffic demands and delay constraints matrices, two parameters (each one with three distinct values) were used, $D_p \in \{0.1, 0.2, 0.3, 0.4\}$ and $DR_p \in \{3, 4, 5\}$, allowing to tune the difficulty levels of traffic demands and delay requirements, respectively (i.e. making the traffic/delay matrices easier or harder to comply). Due to space constraints, we avoid further technical details explanations, and only mention that the precise values of D_p and DR_p were selected in accordance with the overall congestion and delay constraints levels which are intended to be imposed to each network instance. Here, scenarios assuming higher values for the D_p parameter and, simultaneously, lower values of the DR_p parameter are the ones harder to comply. Higher values of D_p mean that higher traffic demands are being considered, being harder to comply. In counterpoint, higher values for the DR_p mean that higher values for edge-to-edge delays requirements are being considered, being easier to comply. Based on the network topology, the demand matrices and a weight assignment to the network links, the framework OSPF simulator will distribute the traffic along the paths, turning possible the computation of the $\Phi^*(w)$ and $\gamma^*(w)$ values. The framework optimization module will then resort to the MOEA optimization approach explained before to find the routing solution for the optimization problem.

2) *Illustrative Results*: As an initial illustrative example of the MOEAs optimization capabilities, a particular optimization scenario was selected of a network with 30 nodes and 110

TABLE II
MAIN INPUTS, OBJECTIVE FUNCTIONS AND OUTPUTS [FIRST OPTIMIZATION SCENARIO]

Inputs	
$G(N, A)$	Representation of the network topology
D	Traffic demands matrix
DR	Delay requirements matrix
Objective Functions	
$\Phi^*(w)$	Congestion measure for the routing solution w
$\gamma^*(w)$	Delay measure for the routing solution w
Output	
$\{w_1, w_2, \dots, w_m\}$	Pareto front with a set of alternative routing solutions

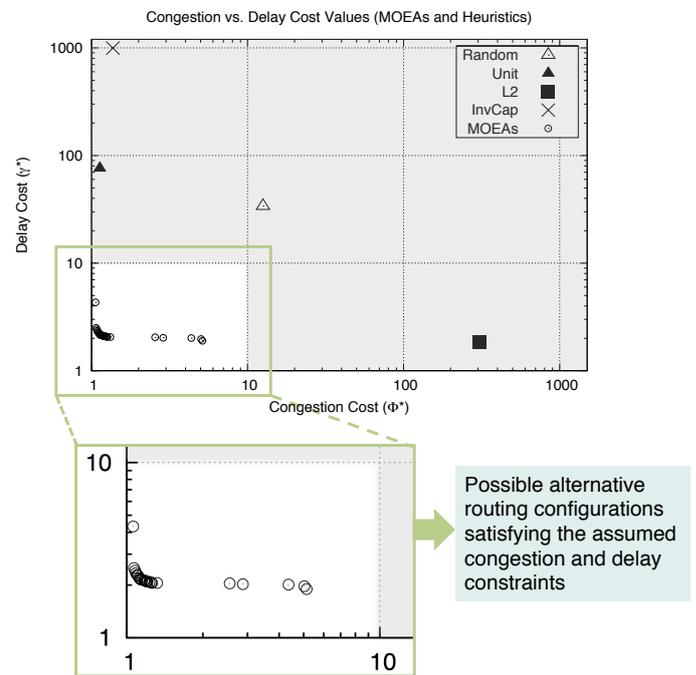


Fig. 4. Comparative perspective of the MOEAs (also zoomed) and heuristics results (logarithmic scale) [scenario with 30 nodes, 110 links, $D_p = 0.1$ and $DR_p = 3$]

links, with $D_p = 0.1$ and $DR_p = 3$. For this particular example, Figure 4 shows a particular subset of the solutions obtained by the MOEAs in the selected scenario and, for comparative terms, the ones obtained by common heuristics described before. Figure 4 has two distinct areas, the first one corresponds to solutions assuming routing configurations able to obey the considered traffic and delay demands (i.e. the white area, where cost function values are lower than 10), and a second area where the routing solutions lead to quality degradation of the network, with overloaded links or with the target delays requests not being assured by the network (gray filled area). Note that in the gray filled at least one of the congestion/delay requirements are not accomplished. Also, a logarithmic scale is used in both axis meaning that most of routing solutions inside the gray filled area represent extremely poor quality routing solutions.

As visible in Figure 4, the solutions obtained by the frame-

work are able to provide the network administrator with a set of near-optimal routing configuration solutions for the network domain (see the Pareto front inside the white area of Figure 4). It is also noticeable that all the results of the heuristics for this instance lie outside the admissible range (outside the white area), and some of them with penalties which are one or two orders of magnitude higher than the ones obtained by MOEAs. This means that none of the heuristics is able to provide acceptable routing configurations for this multi-constrained optimization problem. Figure 4 also magnifies the white area of the figure. As observed, the administrator achieves a set of near-optimal configuration solutions resulting from the Pareto front of the optimization process, all lying inside of the white area. Each one of such points (solutions) is associated with routing weights able to be used by the administrator according with the desired trade-off between the optimization objectives.

Figure 5 shows additional optimization examples taken from two networks with $N \in \{30, 50\}$ nodes, which were selected to also include very complex optimization scenarios. The selected examples assume tree distinct combination of (D_p, DR_p) parameters for each one of the networks, namely $\{(0.1, 3), (0.2, 4), (0.3, 3)\}$. The behavior of the heuristics is similar to the describe before, i.e. they are not capable of achieving acceptable performance, meaning that these solutions are completely outside of the white area of the graphs, not being included in the graphs. Figure 5 results were obtained in the first runs of the MOEA optimization procedures, i.e. only a single run of the optimization algorithm has been performed for each of the network instances. For most of the scenarios, this was sufficient to find acceptable results i.e. Pareto fronts with solutions (i.e. weight settings) in the white area of the figures.

The harder optimization scenarios are the ones imposing higher requirements regarding both the traffic demands and delay restrictions (e.g. $D_p = 0.3, DR_p = 3$). In such scenarios, the Pareto front patterns are not so close to the graph origin as in other network configurations. As observed in Figure 5, the first run of the MOEA was not sufficient to find near-optimal configurations for the last scenarios with $N = 50$ and (D_p, DR_p) values of $(0.3, 3)$. Nevertheless, to improve such preliminary results additional runs of the MOEAs could be used to generate other weight setting configurations overcoming the performance obtained in previous runs. To illustrate the previous reasoning, Figures 6 a), b) plot additional optimization results obtained in other runs of the MOEAs for two specific scenarios. The new Pareto fronts of Figures 6 a), b) are compared with the values obtained in the first runs (observed before in Figure 5). The analysis clearly shows an improvement of the Pareto front patterns for each scenario, containing now several points in the white regions of the figures, i.e. a generalized displacement of the Pareto fronts to the feasible configuration areas of the figures.

3) *Single-objective vs Multi-objective Evolutionary Algorithms*: This section discusses the advantages of using a multi-objective optimization perspective in the devised framework, comparatively with other approaches in the area which assume the same problems under a single objective optimization perspective. As an example, for this specific scenario, a linear

weighting scheme could be used to denote the overall cost of the solution, as expressed by Equation 4.

$$f(w) = \alpha\Phi^*(w) + (1 - \alpha)\gamma^*(w), \alpha \in [0, 1] \quad (4)$$

In Equation 4 parameter (α) determines the importance that is given to each objective (congestion and delays). Examples of performance analysis of this approach can be found in [11], for a large set of distinct scenarios. Nevertheless, this strategy suffers from one main drawback, since it assumes that there is one single trade-off that is optimum. In fact, the algorithms typically return one single solution that has to be implemented by the administrator. To be able to analyze several distinct trade-offs between the two objectives, the user needs to execute different runs of the algorithm using different values of the parameter α . Furthermore, for specific configurations, several tradeoffs between the considered objectives are not possible to be achieved, meaning that it is not straightforward to assign admissible values for the α parameter.

To illustrate this reasoning, in Figure 6 c) we complemented the results presented in Figure 4 with single objective optimization results when run with three distinct values of the parameter α (0.25, 0.5 and 0.75). Figure 6 c) shows a set of selected solutions obtained by the MOEA and also the best solution obtained by each of the single-objective evolutionary algorithms. It is worth also to mention that all the solutions shown for the multi-objective optimization are obtained in a single run, while the solutions for the single objective need three distinct runs. As depicted in Figure 6 c), it is clear that the MOEA approach provides the network administrator with a set of alternatives, from where he can choose the best option, reflecting the ideal trade-off between the two objectives. The single objective evolutionary algorithm (SOEA) is usually restricted to a certain area of the working region of the network, making difficult the tuning process of the importance of each objective.

Thus, the MOEA approach used in the devised framework presents advantages simultaneously regarding the computational efforts perspective and also improving the quality and the diversity of the results provided to the administrator.

B. Resilience to Traffic Variations

The conditions of a networking infrastructure are not static, they rather change over time. As example, network traffic volume suffers several alterations over specific time periods. Although some of the fluctuations on traffic can be unpredictable, others, such as global variations over specific periods of time (e.g. night and day variability) can be foreseen, and translated into distinct estimates by network administrators. Those estimates, represented as traffic matrices, frequently have uncorrelated source-destination individual entries or different overall levels of traffic. In this context, classical TE methods have several difficulties to address such problems, as they usually assume fixed traffic volumes between each source/destination pair of the traffic matrices. As consequence, a routing configuration may be appropriate to warrant a good performance of the network regarding a specific traffic estimation, but can completely fail in respect to another. Finding

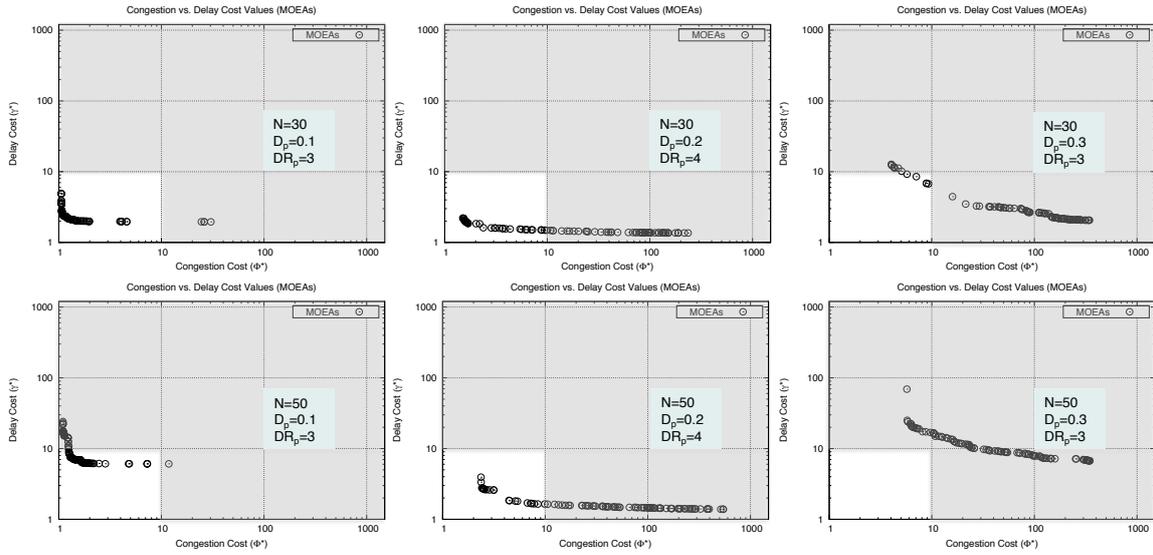


Fig. 5. MOEA optimization results - Pareto fronts - for two synthetic networks ($N = 30, 50$) with distinct combinations of D_p and DR_p parameters (first runs).

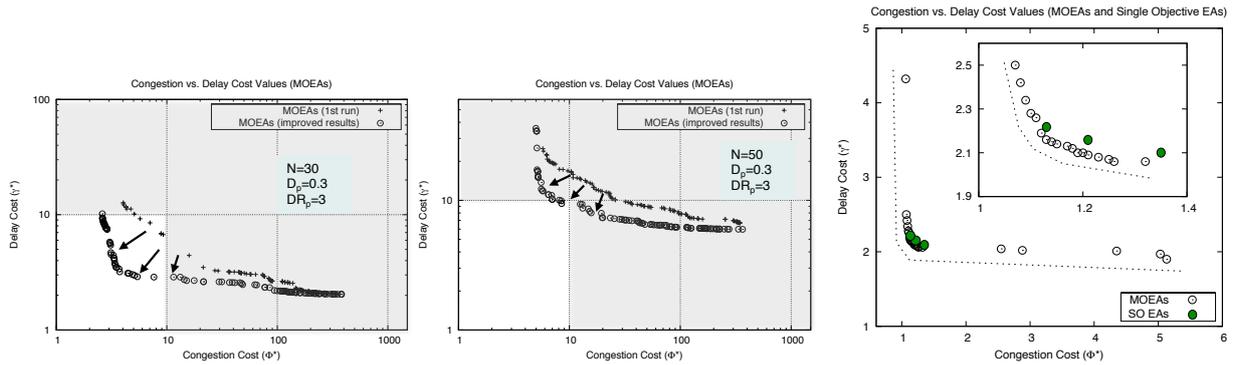


Fig. 6. a)-b) Improved results for three particular network instances of Figure 5 (in runs 2 and 23, respectively); c) Comparison between results from single and multi-objective optimization perspectives.

TABLE III
MAIN INPUTS, OBJECTIVE FUNCTIONS AND OUTPUTS [SECOND OPTIMIZATION SCENARIO]

Inputs	
$G(N, A)$	Representation of the network topology
D_1	First traffic demand matrix
D_2	Second traffic demand matrix
Objective Functions	
$\Phi_1^*(w)$	Congestion measure for the routing solution w under D_1 traffic assumptions
$\Phi_2^*(w)$	Congestion measure for the routing solution w under D_2 traffic assumptions
Output	
$\{w_1, w_2, \dots, w_m\}$	Pareto front with a set of alternative routing solutions

a configuration that is adequate for both can be addressed as a multi-objective problem that is tackled by the proposed framework as follows.

1) *Problem Formulation:* This specific problem formulation assumes, for a given network topology, two traffic demand

matrices expressing distinct traffic aggregates traversing the network infrastructure at distinct time periods (see Table III). Then, the previously explained cost function Φ^* is here used to define two functions, Φ_1^* and Φ_2^* , that evaluate the congestion levels of the network when associated to each of the traffic demand matrices. As consequence, the aim is to find a single configuration of weights (w) that simultaneously minimizes both objectives, Φ_1^* and Φ_2^* , i.e. assures a good distribution of the traffic in both cases.

To test this optimization approach three synthetic network topologies were used, with different sizes (30 and 50 nodes as in the previous optimization scenario) and distinct average in/out degree of each node (2 and 4). For each of these topologies, a set of traffic demands matrices D_i were randomly generated with different levels of traffic amount, in this case with D_p values of (0.3, 0.4, 0.5). The correlation between two matrices in the same scenario is also kept under control with an approximated value 0.5. For comparison purposes, two traditional and commonly used weights configuration schemes are included in the presented results: Unit and InvCap. In addition, to further highlight the benefits of the multi-objective

approach, two single-objective optimizations, that solely minimize one of the two objectives, are also included.

2) *Illustrative Results:* Table IV presents the obtained results for this illustrative scenario. Due to space constraints, and to avoid the presentation of all the Pareto fronts returned by the framework, the results are only the solutions where both objectives are considered by the administrator as equally important. Nevertheless, and as previously explained, it is possible for the administrator to choose from the Pareto solution set a configuration that more accurately reflects its needs (e.g. giving more importance to one of the objectives).

In Table IV, values above the threshold of acceptable congestion on the network are identified with a gray filled background, meaning that in such cases the network was unable to accommodate all the considered traffic demands. The results for the case regarding the optimization for two $D_{0.4}$ traffic matrices on the 30_4 and 50_2 networks are not presented as all values were above the threshold of acceptable congestion cost, meaning that in such scenarios and independently of the optimization method it is not possible to support the considered traffic matrices.

The results show that, in most cases, by resorting to multi-objective optimization, it is possible to obtain a routing configuration that enables the network to perform well even if two disruptive traffic demands matrices are considered. Also presented in Table IV, are the congestion values provided by the traditional schemes, Unit and InvCap, where they totally fail to warrant a good performance level for heterogeneous traffic demands. When the two traffic matrices are divergent, single objective optimizations are also often unable to deliver suitable solutions. Although they provide a good level of congestion for the optimized objective, they do not grant the same level of performance for the other, being the network unable to accommodate the traffic for the unconsidered demand matrix. For most of the cases, the multi-objective optimization algorithm was the only one capable to achieve weight configurations that enable a satisfactory network behavior for the two demands matrices by concurrently minimizing both objectives.

Another important aspect of MOEA is that the optimization mechanisms included in the algorithms, such as the NSGA-II algorithm, allow to attain better congestion levels even in the scope of a single traffic demand matrix weights optimization. This can be observed in the 30_2 topology scenario, with demands level 0.3, where the multi-objective optimization provides a better ranked solution in both objectives. This is mainly due to the diversity of solutions kept within the population during the optimization process.

As a final note it should be highlighted that the optimization approach presented in this section can be adapted to deal with multiple traffic matrices. As obvious, that will make harder the optimization process and, depending on the considered demands matrices, sometimes might be not possible to find a weight solution able to accommodate all the considered demands.

C. Resilience to Link Failures

Link failure events may have severe impact in the performance of a given network infrastructure. As an example, in

link state routing algorithms, when a link fails, the network traffic is shifted to other shortest paths which have meanwhile been recalculated by the routing protocol. However, this relocation of traffic in the network can be problematic, leading to congestion in some parts of it and having serious impact in the offered QoS levels.

Obviously, it is not possible to fully predict link failure events but network administrators often can identify which failure can cause for some reason a significant impact on the network functional conditions (e.g. link capacity, network architecture, failure probability, etc.). Therefore, it would be adequate trying to protect such link against such event, whilst ensuring the continuity of an acceptable congestion level in the network. In this context, this new multi-objective problem aims to minimize the network congestion before and after the failure of a specific link.

1) *Problem Formulation:* The problem can be defined as follows (Table V). For a given network topology, a traffic demand matrix and a previously selected topology link, the target consists in finding a set of weights (w) that minimize simultaneously the objective functions Φ_n^* and Φ_{n-1}^* , which, respectively, assess the congestion cost of the network in a normal state (n), and when the selected link fails ($n-1$). Even in the presence of such link failure, the select routing configuration should be able to assure the quality of the service offered by the network.

2) *Illustrative Results:* The optimization results for the link failure multi-objective problem are shown in Table VI, and, as above, congestion costs obtained by applying traditional weights configurations schemes (Unit and InvCap) and a single objective optimization were also included for comparison purposes. Again, in Table VI, the obtained multi-objective results are those where the solutions give equal importance to both objectives, i.e. we only depicted such unique solution from the Pareto front returned by the framework. The considered single objective optimization only minimizes the congestion for the network normal state, as it would make no sense to optimize the congestion considering only the failing state, disregarding the performance of the network before a link failure that may never take place. In the experiments, the link that belongs to the largest number of shortest paths, when a InvCap weights configuration is applied, was used as the criterion to select the failing link. The developed framework, however, offers a broader set of selection criteria, such as the link with higher load or the link whose failure has the greatest impact on the network congestion cost.

As observed in Table VI, the multi-objective approach was once again the algorithm with best overall results. Even in cases where the threshold is surpassed, such as in the 50_2 topology experiments with demands level 0.4, the MOEA offers a solution set that can assure a near acceptable congestion performance, where all other methods clearly failed, resulting into a network infrastructure with several congestion. It is also important to highlight that sometimes some small penalties can be observed in the congestion cost of the network in the normal state when the multi-objective approach is used. Nevertheless, this is entirely justified by the gain observed in the congestion levels when the network is under link failure. This can be

TABLE IV
CONGESTION OPTIMIZATION FOR TWO TRAFFIC DEMAND MATRICES (MINIMUM VALUES)

Algorithm	Demands #1	Demands #2	30 ₂		30 ₄		50 ₂	
			Φ_1^*	Φ_2^*	Φ_1^*	Φ_2^*	Φ_1^*	Φ_2^*
Unit			83.80	130.68	255.27	160.54	339.96	313.34
InvCap			31.68	15.60	263.04	75.66	437.70	434.55
Single (#1)	0.3	0.3	1.48	5.34	3.63	132.24	2.36	19.88
Single (#2)			6.45	1.56	78.23	2.03	24.91	2.10
Multi-obj			1.45	1.53	3.48	2.22	1.78	1.88
Unit			83.806	227.74	255.27	426.75	339.96	430.14
InvCap			31.68	203.53	263.04	717.95	437.70	812.01
Single (#1)	0.3	0.4	1.48	92.78	4.99	278.04	2.36	84.47
Single (#2)			1.89	2.53	9.50	23.81	13.84	15.05
Multi-obj			1.40	2.51	2.79	17.59	2.05	15.80
Unit			157.00	227.74				
InvCap			221.10	203.53				
Single (#1)	0.4	0.4	2.61	47.54				
Single (#2)			28.91	2.53				
Multi-obj			2.17	2.24				

TABLE V
MAIN INPUTS, OBJECTIVE FUNCTIONS AND OUTPUTS [THIRD OPTIMIZATION SCENARIO]

Inputs	
$G^n(N, A)$	Representation of the network topology (with n links)
$G^{n-1}(N, A)$	Representation of the network topology (with $n - 1$ links)
D	Traffic demand matrix
Objective Functions	
$\Phi_n^*(w)$	Congestion measure for the routing solution w for the topology $G^n(N, A)$
$\Phi_{n-1}^*(w)$	Congestion measure for the routing solution w for the topology $G^{n-1}(N, A)$
Output	
$\{w_1, w_2, \dots, w_m\}$	Pareto front with a set of alternative routing solutions

observed in particular by comparing some single objective optimization results with those provided by the multi-objective optimization approach.

As mentioned before, the values presented in Table VI are the minimum values obtained from the Pareto fronts when equal importance is given to both objectives. However, the administrator is able to select the most appropriate solution from the Pareto front. As an example of a possible choice given to an administrator, in the experiments with the network topology 30₄ and demands $D_{0.4}$, the congestion cost value pair (18.66, 10.13) is presented in Table VI as a representative result. There are, however, other possible solutions in the returned Pareto front, such as the cost value pair (9.22, 34.29) that keeps the congestion on the normal state under the acceptable threshold, but, although better than those provided by the other algorithms, with a worst congestion in a failed state. An administrator should choose which is the most adequate. The developed framework, for that matter, provides a set of tools that can help in the decision making, with several informations regarding individual links usage within each configuration solution. Nevertheless, the results obtained

in all the demand instances and topologies clearly indicate the obvious advantages for an administrator to resort to this preventive multi-objective link failure optimization method.

D. Computational Efforts, Applicability and Availability

The devised framework aims to help network administrators in finding efficient routing configuration for NP-hard optimization scenarios. In this perspective, it is important to assess the computational efforts required to solve such problems. Most of the illustrative optimization scenarios presented in this work involve a considerable number of nodes (i.e. routers), which closely model real OSPF of IS-IS network infrastructures. On these examples, and with the optimization framework running on a current end-user computational platform (e.g. Core i7 processor) a single MOEA's run for the considered examples required a computational time in the order of a few minutes. As obvious, when considering even harder optimization problems, the need of using several MOEAs runs may considerable increase such value. Nevertheless, the overall required computational times are still within acceptable limits considering the utilization context assumed by the proposed framework.

It also should be highlighted that the envisioned application scenarios for the framework are those where the network administrators use this optimization tool to adjust routing configurations for stable and well known time periods. Such optimization efforts are done in conformity with previous studies made in the network domain from which result a set of requirements to be accomplished by the infrastructure (which may express traffic demands requirements, resilience requirements, or any other imposed constraint). Such preliminary optimization of routing configurations allow to attain a more balanced and efficient distribution of the traffic within the network infrastructure, thus having positive cost and performance impact for the Internet Service Providers.

The availability of frameworks and tools like the one presented here is of extreme importance to help network administrators when dealing with complex management and configuration tasks. Thus, a preliminary version of the opti-

TABLE VI
LINK FAILURE CONGESTION OPTIMIZATION (MINIMUM VALUES)

Algorithm	Demands	30 ₂		30 ₄		50 ₂	
		Φ_n^*	Φ_{n-1}^*	Φ_n^*	Φ_{n-1}^*	Φ_n^*	Φ_{n-1}^*
Unit	0.3	130.68	165.30	198.96	234.23	339.96	373.56
InvCap		15.60	88.13	323.76	269.23	437.70	565.62
Single (n)		1.56	30.89	2.00	58.65	1.98	18.09
Multi-obj		1.44	1.48	2.30	2.25	1.77	1.78
Unit	0.4	160.95	165.30	426.75	499.26	339.96	478.33
InvCap		108.28	271.60	717.95	723.39	812.00	919.68
Single (n)		2.03	94.68	7.88	107.13	19.63	97.77
Multi-obj		1.75	1.80	18.66	10.13	11.52	11.49

mization framework described in this paper is made available in the page accessible at <http://darwin.di.uminho.pt/netopt>.

V. CONCLUSIONS

This paper presented an optimization framework for routing configurations based on computational intelligence methods. In particular, Multi-objective Evolutionary Algorithms are used to solve complex optimization problems pursuing near-optimal network configurations able to improve the resilience levels of network infrastructures. As a proof-of-concept, three illustrative TE methods were described along with illustrative optimization results. First, a multi-constrained scenario was assumed, involving traffic and delay constraints imposed by the network administrator. The second method allows to achieve network routing configurations that are robust to changes in the traffic demands traversing the infrastructure, which are expressed by traffic matrices. The third proposed mechanism ensures that the network continues to operate with an appropriate level of quality even in the presence of fault situations of certain infrastructure links. In all the cases, the network operator is able to select a specific solution, from a computed Pareto front, representing the most appropriate trade-off between the considered objectives.

The presented optimization results clearly corroborate the effectiveness of the used optimization engines on solving complex network optimization problems. Thus, the devised optimization framework is a valuable tool for network administrator allowing for automated optimization of network infrastructures. In particular, the multi-objective nature of the outputs provided by the framework allow that network administrators select the solutions with the most appropriate tradeoff between the considered objectives.

ACKNOWLEDGMENTS

This work has been supported by COMPETE: POCI-01-0145-FEDER-007043 and FCT Fundação para a Ciência e Tecnologia within the Project Scope: UID/CEC/00319/2013.

REFERENCES

- [1] Lee, K., Lim, F., Ong, B.: Building Resilient IP Networks. Cisco Press (2012)
- [2] J. Moy. OSPF Version 2. RFC 2328 (Standard), April 1998. Updated by RFC 5709.
- [3] H. Gredler and W. Goralski. The Complete IS-IS Routing Protocol, Springer, 2005.
- [4] Cisco Customer Case Study, Maximizing Return on Network Infrastructure Investment with Cisco MATE Design, 2013.
- [5] M. Suchara, D. Xu, R. Doverspike, D. Johnson, J. Rexford. Network Architecture for Joint Failure Recovery and Traffic Engineering, Proceedings of SIGMETRICS11 Conference, 2011.
- [6] Aysegil Altin, Bernard Fortz, Mikkel Thorup, and Hakan Umit. Intra-domain traffic engineering with shortest path routing protocols. Annals of Operations Research, 204(1):56-95, 2013.
- [7] B. Fortz and M. Thorup. Optimizing ospf/is-is weights in a changing world. IEEE Journal on Selected Areas in Communications, 20(4):756-767, 2002.
- [8] T. Szigeti, C. Hattingh, R. Barton, K. Briley. End-to-End QoS Network Design: Quality of Service for Rich-Media & Cloud Networks, Cisco Press, 2nd Edition, 2013.
- [9] Cariden Technologies. Building Traffic Matrices: Introduction to MATE Flow Collection. White Paper - Version 2. (October 2012)
- [10] M. Rocha, P. Sousa, P. Cortez, and M. Rio. Quality of Service Constrained Routing Optimization Using Evolutionary Computation. Applied Soft Computing, 11(1):356364, 2011.
- [11] Pedro Sousa, Miguel Rocha, Miguel Rio, Paulo Cortez, Efficient OSPF Weight Allocation for Intra-domain QoS Optimization, IPOM 2006 - 6th IEEE International Workshop on IP Operations and Management. Dublin, Ireland, LNCS 4268, pp. 37-48, Ed. Gerard Parr, David Malone and Mchel Foghl, Springer-Verlag, October 2006.
- [12] Vitor Pereira, Miguel Rocha, Paulo Cortez, Miguel Rio, Pedro Sousa, A Framework for Robust Traffic Engineering using Evolutionary Computation, 7th International Conference on Autonomous Infrastructure, Management and Security (AIMS 2013), Barcelona, Spain, Springer, LNCS 7943, pp. 2-13, 2013.
- [13] Pedro Sousa, Miguel Rocha, Miguel Rio, Paulo Cortez, Automatic Provisioning of QoS Aware OSPF configurations, Journal of Networks (JNW), 2(2):1-10, Academy Publisher, April 2007. ISSN: 1796-2056.
- [14] M. Ericsson, M. Resende and P. Pardalos. A genetic algorithm for the weight setting problem in OSPF routing. Journal of Combinatorial Optimization, 6, 299-333, 2002.
- [15] B. Fortz. Internet traffic engineering by optimizing ospf weights. In Proceedings of IEEE INFOCOM, pages 519-528, 2000.
- [16] W. Ben-Ameur, E. Gourdin, B. Liau and N. Michel. Optimizing administrative weights for efficient single-path routing. In Proc. of networks, 2000.
- [17] A. Bley, B. Fortz, E. Gourdin, K. Holmberg, O. Klopfenstein, M. Piro, A. Tomaszewski, and H. Ümit. Optimization of OSPF routing in IP networks. In A. M. C. A. Koster & X. Muoz (Eds.), Graphs and algorithms in communication networks: studies in broadband, optical, wireless and ad hoc networks (pp. 199-240). Berlin: Springer. Chap. 8, 2010.
- [18] D. Xu, M. Chiang, and J. Rexford. DEFT: Distributed exponentially-weighted flow splitting. Proc. IEEE Conf. Comput. Commun., pp. 71-79, 2007.
- [19] V. Pereira, M. Rocha and P. Sousa, Optimizing Load Balancing Routing Mechanisms with Evolutionary Computation. Proc. International Workshop on Future Internet and Smart Networks (FI&SN 2016), IOS Press, Ambient Intelligence and Smart Environments Volume 21, pp. 298-307, Intelligent Environments, 2016.
- [20] A. Altin, B. Fortz, M. Thorup, and H. Ümit. Intra-domain traffic engineering with shortest path routing protocols. Annals of Operations Research, 204, issue 1, pp. 65-95, Springer, 2013

- [21] Tune, P., Roughan, M.: Network design Sensitivity Analysis. In The 2014 ACM International Conference on Measurement and Modeling of Computer Systems. ACM, SIGMETRICS 14, 449-461 (2014)
- [22] Dahai Xu, Mung Chiang, and Jennifer Rexford. 2011. Link-state routing with hop-by-hop forwarding can achieve optimal traffic engineering. IEEE/ACM Trans. Netw. 19, 6 (December 2011)
- [23] E. Dijkstra. A note on two problems in connexion with graphs. Numerische Mathematik, 1(1):269-271, 1959.
- [24] E. Zitzler, M. Laumanns, and L. Thiele. Spea2: Improving the strength pareto evolutionary algorithm. Technical report, 2001.
- [25] Kalyanmoy Deb, Samir Agrawal, Amrit Pratap, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. IEEE Trans. Evolutionary Computation, 6(2):182-197, 2002.
- [26] B. Fortz. Internet Traffic Engineering by Optimizing OSPF Weights. In Proceedings of IEEE INFOCOM, pages 519-528, 2000.
- [27] A. L. A. Medina, I. Matta, and J. Byers, BRITTE: Universal Topology Generation from a Users Perspective, Tech. Rep. 2001-003, Jan. 2001.
- [28] Coello, C.: A Comprehensive Survey of Evolutionary-Based Multiobjective Optimization Techniques. Knowledge and Information Systems 1(3), 129-156 (1999)



Pedro Sousa graduated in Systems and Informatics Engineering at the University of Minho, Portugal, in 1995. He obtained a MSc Degree (1997) and a PhD Degree (2005), both in Computer Science, at the same University. In 1996, he joined the Computer Communications Group of the Department of Informatics at University of Minho. He is an Assistant Professor and performs research activities within Centro ALGORITMI at the same university, mainly focusing on Computer Networks related topics (see <http://marco.uminho.pt/~pns> for more information).



Vítor Pereira received the Licentiate degree in Mathematics in 1998, the Licentiate degree in Computer Sciences in 2008, and the MSc in Networks and Communication Services in 2012 from the University of Minho, Portugal. He is currently a PhD candidate in Informatics, and a researcher at the ALGORITMI Center at the University of Minho. His research interests include Internet design, control, and management; nonlinear network optimization and Evolutionary Multi-Objective Optimization.



Paulo Cortez (Habilitation, PhD) is an Associate Professor at University of Minho and Coordinator of the Information Systems and Technologies R&D group of ALGORITMI Research Centre. His research interests include: intelligent data analysis, forecasting and optimization. Currently, he is an associate editor of the Expert Systems journal. His research has appeared in Journal of Heuristics, Decision Support Systems and others (see <http://www3.dsi.uminho.pt/pcortez>).



network services.

Miguel Rio received the Ph.D. degree from the University of Kent, Canterbury, U.K., and the M.Sc. and M.Eng. degrees in informatics from the University of Minho, Braga, Portugal. He is a Reader (Associate Professor) of Computer Networks, Department of Electronic and Electrical Engineering, University College London, London, U.K. He has authored extensively in top ranked conferences and journals. His research interests include network measurement, congestion control, new network architectures, and, more recently, in the interaction between cloud and



Miguel Rocha is currently an Associate Professor at the Informatics Department, University of Minho, Portugal, the Director of the MSc in Bioinformatics and a senior researcher within the Centre of Biological Engineering (CEB) where he co-leads a research team in Bioinformatics and Systems Biology. He is the author of around 150 publications in international journals and conferences. Also, he has been the PI and has collaborated in several funded research projects by the Portuguese FCT, European Commission and private companies.