Multiobjective Genetic Algorithms for Solving the Impairment-Aware Routing and Wavelength Assignment Problem

Demetris Monoyios and Kyriakos Vlachos

***Abstract*—In future transparent (all-optical) WDM networks, the signal quality of transmission (QoT) will degrade due to physical layer impairments. In this paper, we propose two genetic algorithms for solving the static impairment-aware RWA (IA-RWA) problem by accounting for the impact of physical impairments in the optimization process when searching for the optimum routing path and wavelength channel. The first algorithm indirectly considers the physical impairments through the insertion of the path length and the number of common hops in the search process, using classical multiobjective optimization (MOO) strategies. The second algorithm is a single-objective genetic algorithm (GA) that uses the Q factor for the evaluation of the feasibility of the selected RWA solution. The Q factor is used in each iteration of the algorithm in a self-learning mode in order to evaluate the fitness of each solution to the RWA problem and trigger the evolution of the population. Performance results have shown that considering path length and number of common hops for indirectly handling impairments provide an efficient solution to the IA-RWA problem.**

***Index Terms*—RWA; Genetic algorithm; All-optical networks; Multiobjective optimization**.

I. INTRODUCTION

Introducing optical transparency in the optical layer on one hand leads to a dynamic, flexible optical layer with the possibility to add extra "intelligence." On the other hand, transparency reduces the ability of the digital electronic layer to interact with the optical layer. Thus, optical transparency has an impact on network design, either by adapting the size of WDM transparent domains in order to neglect the physical impact on the quality of transmission (QoT) or by introducing physical considerations in the network and wavelength planning process. The problem of setting up lightpaths by routing and assigning a wavelength to each connection is called the *routing and wavelength assignment* (RWA) problem [1]. The RWA problem belongs to the category of NP-complete problems; that is, the computational time would increase exponentially with the problem size. Thus, a wide range of optimization methods and heuristics have been proposed to solve various optical network optimization problems, as, for example, integer linear programming [1]. Their high complexity, however, requires intensive computational efforts, and therefore other novel heuristic algorithms such as Tabu search [2], simulated annealing [3], and genetic algorithms [4,5] have been used.

Genetic algorithms are stochastic search optimization methods that are widely used in combinatorial optimization and parameter tuning applications. They have been used for solving the plain single-objective RWA problem [4], for optimizing amplifier placement, and for multicasting sessions [5]. When, however, fiber impairments are considered, most of these studies consider the QoT problem separately from the RWA problem [6]. If the Q factor is greater than the requested quality, the connection is accepted, or otherwise it is dropped and a new candidate path is calculated. Another approach is to incorporate impairments into the cost function of the RWA problem. However, a cost function that considers both linear and nonlinear impairments is still an open issue. Therefore, other metrics have been used, including the average measured Q [7] or the noise variance [8]. In any case, accurate Q-path estimation is a heavy computational task, and even in the static RWA problem, with offline processing, it requires calculations for all combinations of the k-shortest paths for all source–destination pairs.

In this paper, we propose the use of two genetic algorithms for solving the IA-RWA problem. The first is a multiobjective genetic algorithm (MOGA) that uses classical multiobjective optimization strategies to jointly solve the IA-RWA problem. This algorithm decomposes the problem into its routing (R) and wavelength assignment (WA) subproblems. It then uses path length and number of common hops, inserted as entries in a multiobjective vector, that must be optimized jointly. This allows for the calculation of a global optimum solution that is a set of routes that can serve the requested connections with the minimum number of wavelengths and also have an acceptable QoT performance. The second IA-RWA algorithm is a single objective, Q-learning scheme. The algorithm calculates, during each iteration, the Q factor of each lightpath before evolving the genetic population to the next generation. Such an approach is computa-

Manuscript received November 2, 2009; revised August 24, 2010; accepted September 9, 2010; published December 15, 2010 (Doc. ID 118639).

The authors are with the Computer Engineering and Informatics Department and Research Academic Computer Technology Institute, University of Patras, Rio, Greece (e-mail: kvlachos@ceid.upatras.gr).

Digital Object Identifier 10.1364/JOCN.3.000040

tionally hard, but is useful for both solving the offline IW-RWA problem as well as for evaluating heuristic approaches like the first proposed one.

The rest of the paper is organized as follows. Section II presents the MOGA for solving the IA-RWA problem, while Section III presents the single-objective, Q-learning genetic algorithm. Section IV presents and compares evaluation results of both approaches. Finally, Section V concludes the paper.

II. MULTIOBJECTIVE GENETIC ALGORITHM FOR SOLVING THE IA-RWA PROBLEM

In this section, we adapt multiobjective optimization techniques for solving the IA-RWA problem. To do so, we formulate the problem and introduce novel fitness functions in order to indirectly account for the effect of the physical impairments on the QoT of each lightpath.

A. Problem Definition

We are given a network represented by a connected graph $G=(N,E)$, where N denotes the set of nodes (not equipped with wavelength conversion capabilities), and E denotes the set of point-to-point single-fiber links. Each fiber link $l \in E$ is able to support W distinct wavelengths. We are also given a traffic matrix Λ of nonnegative integers, where Λ_{sd} denotes the number of connections that have to be established from source node s to destination node d. Each connection requires bandwidth equal to one wavelength, and there may be multiple connection requests for a given source–destination pair (s,d). Each established connection is categorized as feasible or not based on its QoT. Finally, we let W be the number of distinct wavelengths that have been used to serve all connection requests. The aim of the IA-RWA algorithm is to minimize the number of used wavelengths W, while selecting lightpaths that have acceptable QoT performance for servicing connection requests.

B. Basic Formulations

The multiobjective genetic algorithm proposed in this paper extends the genetic algorithm (GA) presented in [4]. In particular, a gene corresponds to one of the k-shortest paths of the P_{sd} set that serves a specific source–destination pair (s,d). To this end, a set of $\Sigma_{(s,d)pairs}\Lambda_{sd}$ randomly generated genes comprises a chromosome, which constitutes a potential solution to the routing problem of the IA-RWA instance.

The objective of the genetic algorithm is to find those genes that minimize the number of wavelengths needed for coloring the resulting chromosome. Basic functions for doing so starting from a randomly initiated population and performing iterations are to rank and select chromosomes based on some fitness function, perform mutation/crossovers to avoid local optimums, and finally select the best-performing ones. In what follows, we introduce novel fitness functions that will define the objective vectors for the multiobjective genetic algorithm, basic genetic operands, and also present a wavelength assignment heuristic for assigning wavelengths to the selected best-performing chromosomes.

Fitness functions: In order to reach a solution with a minimum number of wavelengths used, it is important to avoid the repeatability of the links in the paths (genes) that comprise the chromosome. Thus, we denote by F_l the cost of a link l, and we assume that is equal to the number of paths that cross link l. To this end, in our formulation, we assign to each gene (s,d) of the chromosome the following cost:

$$c_{sd} = \sum_{l \in p} F_l, \quad p \in P_{sd} \quad \text{and} \quad p \in Ch.$$

This cost directly relates the repeatability of a link in a chromosome, which in turn determines the number of wavelengths to be used. We then define the following fitness function of the chromosome Ch:

$F(Ch)=N^{\mu(c_{sd})}$, where $\mu(.)$ is the average of the cost values that comprise Ch. It has been shown [9] that such a fitness function performs better, when searching for a global optimum solution, than exponential costs [4].

In order to define those fitness functions that will consider the QoT during the multiobjective optimization process, we have categorized the physical impairments according to how they affect the QoT. First, amplified spontaneous emission noise (ASE) depends on the number of erbium-doped fiber amplifiers (EDFAs) of the path, which in turn depends mainly on the length of the path. Polarization mode dispersion (PMD) and chromatic dispersion also depend on the length of the path. Other impairments, such as crosstalk (XT) and filter concatenation (FC), have a significant effect when two (or more) signals go through common cross-connects (OXCs). Thus the number of common hops across the routing path can be a relative objective function to optimize. Driven by these observations, we assign to each chromosome a vector F of objective functions to optimize (minimize) the number of wavelengths as well as the path length and number of common hops for all lightpath requests, in order to meet their QoT restrictions.

For a single gene that corresponds to a source–destination pair (s,d), we assign a cost c_{sd} that denotes the utilization of links that comprise the path $p \in P_{sd}$ and a length cost $d_{sd}=d_p$, where d_p is the length of the selected path $p \in P_{sd}$ and $p \in Ch$. In addition, we also define F_n as the number of paths in that chromosome that utilize node $n \in V$. Thus, we can assign to each gene (s,d) an additional cost h_{sd} that reflects the number of common hops between the selected path for that (s,d) pair and all the other selected paths that comprise the chromosome as follows:

$$h_{sd} = \sum_{n \in p} F_n, \quad p \in P_{sd} \quad \text{and} \quad p \in Ch.$$

To this end, based on the above formulations, we assign to a chromosome Ch the following objective vector:

$$F(Ch) = \{f_1(Ch), f_2(Ch), f_3(Ch)\},$$

where

i. $f_1(Ch)=N^{\mu(d_{sd})}$ is the length cost,
ii. $f_2(Ch)=N^{\mu(h_{sd})}$ is the common hop cost, and

iii. $f_3(Ch) = N^{\mu(c_{sd})}$ is the link utilization cost.

The parameters d_{sd}, h_{sd}, and c_{sd} are as defined above, and $\mu(.)$ denotes the average of the corresponding cost values of the genes inside the chromosome Ch.

Domination relation: Between two chromosomes with two different objective vectors denoted as $F(Ch_1) = (f_1(Ch_1), f_2(Ch_1), f_3(Ch_1))$ and $F(Ch_2) = (f_1(Ch_2), f_2(Ch_2), f_3(Ch_2))$, $F(Ch_1)$ dominates $F(Ch_2)$ if $(\forall f_i(Ch_1) \leq f_i(Ch_2))$ and there exists at least an f_i for which it holds $f_i(Ch_1) < f_i(Ch_2)$; see Fonseca and Flemming [10].

In principle, the three functions of the vector do not contribute equivalently and this depends on the actual network under study. However, this must not be confused with excluding one function from the vector or adding weighting factors. This is part of the optimization process and the result (the finally selected individuals) will clearly be a set of nondominated individuals.

Genetic operands: A n-point crossover and a random uniform mutation with a probability of $1/\Sigma_{(s,d)pairs}\Lambda_{sd}$ are used.

Wavelength assignment heuristic: In the proposed, multiobjective genetic algorithm, we use a wavelength assignment heuristic algorithm, denoted here as WA(Ch), for coloring the best-performing chromosome. This heuristic algorithm is an extension of the Brelaz heuristic [11]. The prime idea is to avoid crosstalk-generating sources [6], so as to obtain solutions that have better quality of transmission performance. To this end, we construct the common edge graph and the common node graph of the genes that comprise the chromosome. By definition, the common node graph has at least the same edges as the common edge graph, since two paths that have a common link also have two common nodes. Typically, a wavelength assignment algorithm is performed on the common edge graph, but in order to avoid interchannel crosstalk, the graph of common nodes has to be used instead. The interchannel crosstalk is the effect of power leakage between lightpaths crossing the same switch and using the same wavelength due to nonideal isolation of the inputs/outputs of the switching fabric. Note that interchannel crosstalk has an imminent effect on QoT because it cannot be filtered out since the interfering signal is on the same wavelength as the one affected.

Moreover, the heuristic WA(Ch) algorithm uses the definition of the number of common hops cost parameter h_{sd} that was also used in one of the fitness (objective) functions. The cost h_{sd} is equal to the number of common hops between the selected path for that (s,d) gene and all the other selected paths that comprise the chromosome. If $\bar{h}$ is the average number of common hops over all the selected paths in the chromosome, our heuristic algorithm iteratively colors each one of the genes (s,d) as follows:

If $h_{sd} < \bar{h}$, then *do the wavelength assignment based on the common edge graph, else do the wavelength assignment based on the node graph.*

Thus, if the number of common hops of the path that corresponds to (s,d) is smaller than the average number of common hops $\bar{h}$, we perform a typical wavelength assignment process on the common edge graph, so as to satisfy the distinct wavelength assignment constraint. On the other hand, if the number of common hops of the examined path is larger than $\bar{h}$, meaning that many other paths utilize the same nodes and the number of crosstalk-generating sources is high, we avoid using the same wavelength with the other ligthpaths by performing the wavelength assignment process on the common node graph. Note that since the common node graph includes the edges of the common edge graph, the distinct wavelength assignment constraint is also satisfied. In this way, the number of the crosstalk-generating sources is decreased, but, however, with a small increase in the required number of wavelengths to color the chromosome.

III. Q-LEARNING GENETIC ALGORITHM FOR SOLVING THE IA-RWA PROBLEM

In this section, we present our second genetic algorithm for solving the IA-RWA problem. This algorithm uses the Q factor to evaluate the fitness of the chromosomes. Such a Q-learning scheme can be considered as the ideal indicator for population evolution towards a solution that consists of only feasible lightpaths. Therefore, we use a single-objective genetic algorithm that is connected to a VPI photonic simulator [12] to estimate the Q factor. The pseudocode of the algorithm is presented in Fig. 1. The optimization criterion is the acceptable QoT for all the connection requests. If this criterion cannot be satisfied, then the optimization criterion is changed to a maximum number of elapsed epoch loops (not shown in the pseudocode of Fig. 1), depending on the amount of the population of the GA, unless progress in the evolution is observed. In such a case the optimization criterion is set back to the acceptable QoT for all the connection requests.

Again, each gene represents one solution for a specific

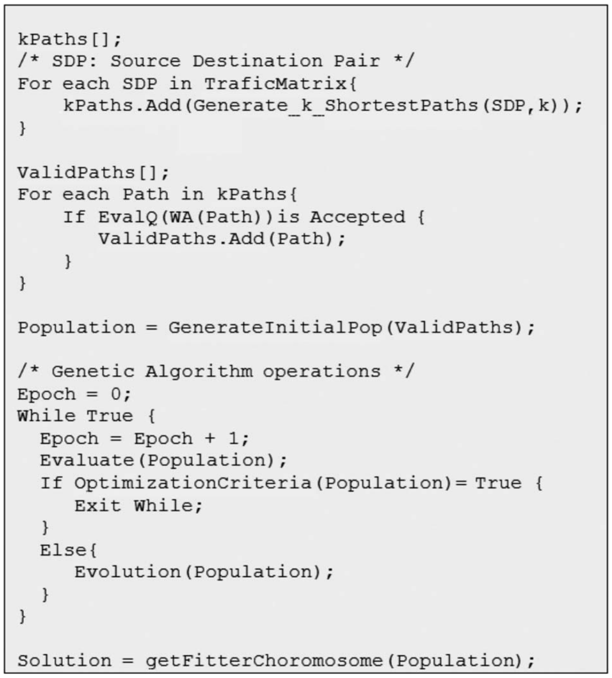

```
kPaths[];
/* SDP: Source Destination Pair */
For each SDP in TraficMatrix{
    kPaths.Add(Generate_k_ShortestPaths(SDP,k));
}

ValidPaths[];
For each Path in kPaths{
    If EvalQ(WA(Path))is Accepted {
       ValidPaths.Add(Path);
    }
}

Population = GenerateInitialPop(ValidPaths);

/* Genetic Algorithm operations */
Epoch = 0;
While True {
  Epoch = Epoch + 1;
  Evaluate(Population);
  If OptimizationCriteria(Population)= True {
     Exit While;
  }
  Else{
     Evolution(Population);
  }
}

Solution = getFitterChoromosome(Population);
```

Fig. 1. Pseudocode of the Q-learning GA algorithm.

source–destination pair, whereas a chromosome represents a solution to the routing subproblem of the RWA instance. Note that we feed to the photonic simulator the candidate path with the highest possible QoT, whereas a path that, even if examined alone, does not have acceptable transmission quality is immediately dropped. This is because the presence of other lightpaths would nevertheless deteriorate its QoT further. In this way, we initialize the population to consist of paths that yield feasible solutions.

Fitness function: Since a chromosome Ch corresponds to a candidate solution of the routing subproblem of the RWA instance and in order to use the physical layer evaluation module, we need first to assign wavelengths. Let $\mathrm{WA}(Ch)$ be a coloring solution to the routes of chromosome Ch. Let also $\mathrm{Eval}(\mathrm{WA}(Ch))$ be the blocking performance of chromosome Ch. Assuming that the $\mathrm{WA}(Ch)$ algorithm has reached a solution with W wavelengths, we define the following fitness function F for the chromosome Ch:

$$F(Ch) = B\ \mathrm{Eval}(\mathrm{WA}(Ch)) + W,$$

where B is a large integer constant number (e.g., $B = 1000$). In this way, the fitness of a chromosome depends on a weighted combination of blocking and the number of used wavelengths. Emphasis is given to the blocking performance by selecting large values for the B coefficient. When blocking is zero, and thus $\mathrm{Eval}(\mathrm{WA}(Ch)) = 0$, then the fitness of the chromosome depends only on the number of wavelengths W that the WA(.) outputs. To this end, we may assume that the initial population consists of chromosomes with nonzero physical layer blocking, and the algorithm tries to reach solutions with zero blocking, without examining to a great extent the number of wavelengths that they use. Once such solutions have been reached, the algorithm tries to minimize the number of used wavelengths. Crossover and mutation operators are also applied, as in the previous MOGA scheme.

Fig. 2. (Color online) DTAG/T-Systems National Core Network.

TABLE I
SIMULATION PARAMETERS

k-shortest paths	2
Initial population	50
Max population	75
Epochs (for single objective)	300
Q threshold	15
Po	−4 dBm
EDFA NF	6 dB
Span	80 km
Channel Spacing	100 GHz

IV. PERFORMANCE EVALUATION

A. Blocking Performance of the MOGA Scheme

We have evaluated the performance of the proposed multiobjective and single-objective genetic algorithms considering the network topology of Deutsche Telecom shown in Fig. 2. The parameters of the simulations are summarized in Table I.

The span length in each link was set to 80 km. Each link was assumed to consist exclusively of standard single-mode fiber (SSMF) with dispersion parameter $D = 17$ ps/nm/km and attenuation parameter $a = 0.25$ dB/km. For the dispersion-compensating fiber (DCF) we assumed parameters $a = 0.5$ dB/km and $D = -80$ ps/nm/km. The launch power was 3 dBm/ch for every SMF span and −4 dBm/ch for the DCF modules. The EDFAs' noise figure was 6 dB, and each EDFA exactly compensates for the losses of the preceding fiber span. Regarding the dispersion management scheme, a precompensation module was used to achieve better transmission reach.

Figure 3 displays the blocking probabilities obtained in our experiments as a function of connection requests (source–destination pairs). In particular, Fig. 3 displays the performance of the MOGA scheme when i) optimizing the RWA problem against fiber impairments [see the MOGA1 curve optimizing only f_1 (path length) and f_2 (common hops)], ii) optimizing fiber impairments as well as the

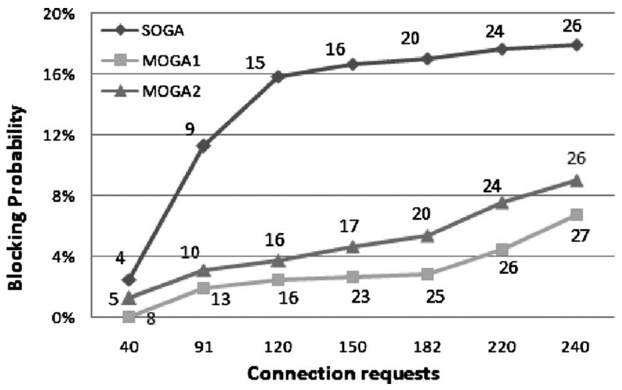

Fig. 3. Blocking probability versus load (number of active source–destination pairs). The tags next to the points indicate the number of wavelengths required.

number of wavelengths (see the MOGA2 curve that optimizes f_1 and f_2 and minimizes link utilization cost f_3), and iii) optimizing only link utilization cost f_3 (single-objective genetic algorithm; see the SOGA curve).

The latter actually corresponds to the pure RWA problem without considering fiber impairments and has been added for reference purposes. Tags in Fig. 3 denote the number of wavelengths required to serve the connection requests. From Fig. 3, it is clear that when considering fiber impairments in the RWA, blocking decreases significantly. This is because many solutions found by the SOGA are not considered by MOGA schemes as accepted since they do not satisfy the QoT.

As Fig. 3 shows, blocking has been decreased from 18% (for connection requests larger than 220) down to 5% (see the MOGA1 and SOGA curves). This indicates that by choosing the set of paths that exhibit short lengths and also have a small number of common hops, one can increase the quality of transmission of the selected lightpaths. However, this enhancement in blocking has resulted in an increase in the number of wavelengths required to serve all connections with the requested QoT. This is because more wavelengths are needed in order to avoid the degradations coming from the physical impairments. However, when the link utilization cost f_3 is also inserted into the optimization process (see the MOGA2 curve), the required number of wavelengths is again decreased. Indeed, from Fig. 3, it can be seen that the number of wavelengths used by MOGA2 is smaller than that used by MOGA1, and for loads higher than 150 requests the number of wavelengths used is the same as that of the SOGA. The resulting trade-off, however, is that the blocking performance of MOGA2 has slightly increased. This is an inevitable trade-off, since the use of fewer wavelengths means that wavelengths are more densely placed, and thus effects like XT, XPM, and FWM have a high effect in signal QoT.

B. Convergence Performance of the MOGA Scheme

In evolutionary algorithms, like genetic algorithms, it is necessary to find whether the algorithm converges and how many epochs are needed. Here, we make use of the Pareto converging genetic algorithm [13] to gain knowledge of the propagation of the solution front through successive generations of the population. This can serve as a clue of convergence speed and convergence performance and will further explore the Pareto front.

Figure 4 displays the intrapopulation ranking histogram for 40 connection requests. This histogram actually displays the ratio of a given rank in the current population to that of the combined and reranked population of the current and the preceding epoch. Therefore, convergence is achieved when, between two successive evolutions, individuals remain nondominated. In that case, the rank ratio is 0.5. To this end, from Fig. 4, it can be seen that the MOGA scheme converges at epoch 357 (mean execution time 5 ms/epoch). Beyond 357 epochs, all subsequent generations consist of only nondominated chromosomes, which is an indication that optimization has converged.

Another histogram that has been used for exploring the Pareto front and global convergence performance is the interpopulation ranking histogram. This combines and reranks different final populations (potentially optimal Pareto solutions) in order to demonstrate whether these individuals are close (or not) to the true Pareto front. It can be the case that no evolutionary gain is achieved beyond a certain number of population evolutions. Figure 5 displays the corresponding interrank histogram between two populations for different initial populations and for different numbers of epochs. It is clear that by epoch 357 the solution found for an initial/maximum population size of 50/75 is not close to the Pareto front since some individuals have been degenerated and are now dominated by some others (ranks 5, 10, and 35; see blue bars in Fig. 5). Further, it was found that there was no further gain from increasing the number of epochs. This is due to the limited search space (small ini-

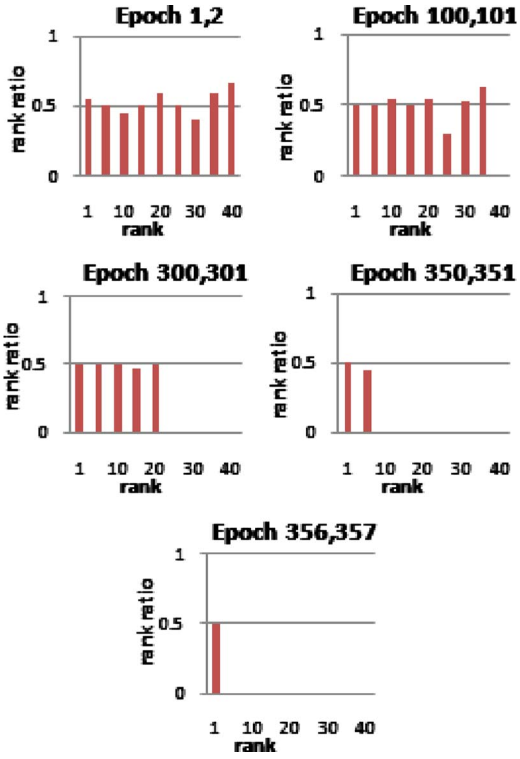

Fig. 4. (Color online) Intrapopulation ranking histogram displaying the ratio of a certain rank of genes between two successive generations.

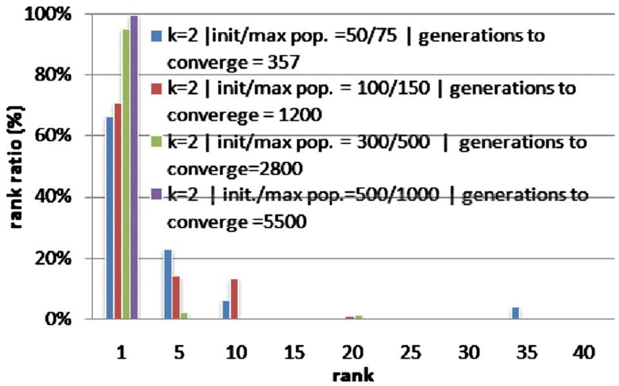

Fig. 5. (Color online) Interrank histogram of two merged optimal solutions on several experiments with different parameters.

tial population), and therefore, we have experimented with larger populations and a number of epochs. As expected, (see the red, green, and purple bars in Fig. 5), better results were obtained by increasing the initial as well as the maximum population limit since the algorithm searches in a wider set of solutions. However, this had a significant cost in terms of computational time since a significantly higher number of iterations were needed to achieve convergence. From Fig. 5, it can be seen that total convergence (rank ratio=1) has been achieved in the case of having 500 individuals for the initial population and 1000 as the maximum. In such a case, the MOGA converges to the optimal Pareto front by the 5500th epoch. This is important information, since being on (or even close) to the true Pareto front is a good indication of the reliability of the proposed genetic algorithms for solving the IA-RWA problem.

Other scenarios with more than 40 connections have also been simulated. In all the cases, solutions have reached a Pareto front (according to intraranking histograms) and in most cases a true Pareto front (according to the interranking histogram).

Figure 6(a) and 6(b) displays a representative evolution of the population in the case of 40 connection requests and an initial/maximum population size of 500/1000. In particular, Fig. 6(a) shows the evolution of the population from the 0th to the 357th epoch, whereas Fig. 6(b) shows a detailed snapshot close to final convergence (from the 3000th to the 5500th epoch). In Fig. 6(b), the (potential) optimal Pareto front has been drawn. From Fig. 6, it is clear that the case of having an initial/maximum population of 500/1000 does converge to a solution that potentially is close to (or on) the Pareto front. The blocking performance of this case is also improved (see Fig. 7) and is discussed in the next subsection.

C. Performance Evaluation of the Q-Learning Genetic Algorithm

In order to evaluate the single-objective, Q-learning genetic algorithm, we carried out simulation experiments again. The simulation parameters remained the same as in Table I. Figure 7 displays and compares the performance of the Q-learning with the multiobjective genetic algorithm of different initial/maximum populations. In particular, Fig.

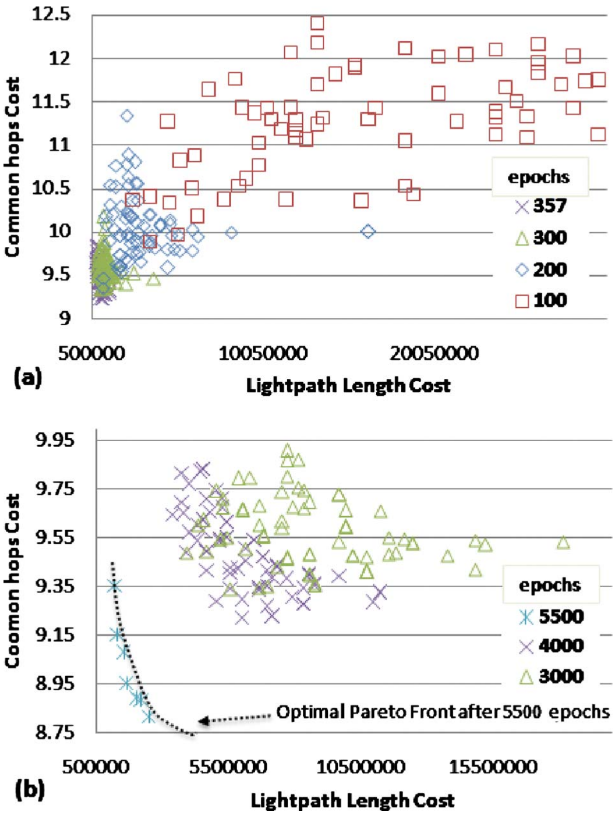

Fig. 6. (Color online) (a) Population evolution towards convergence against the number of hops and the path length for the MOGA1 algorithm for the DTnet and 60 lightpath requests. (b) The converged optimal Pareto front.

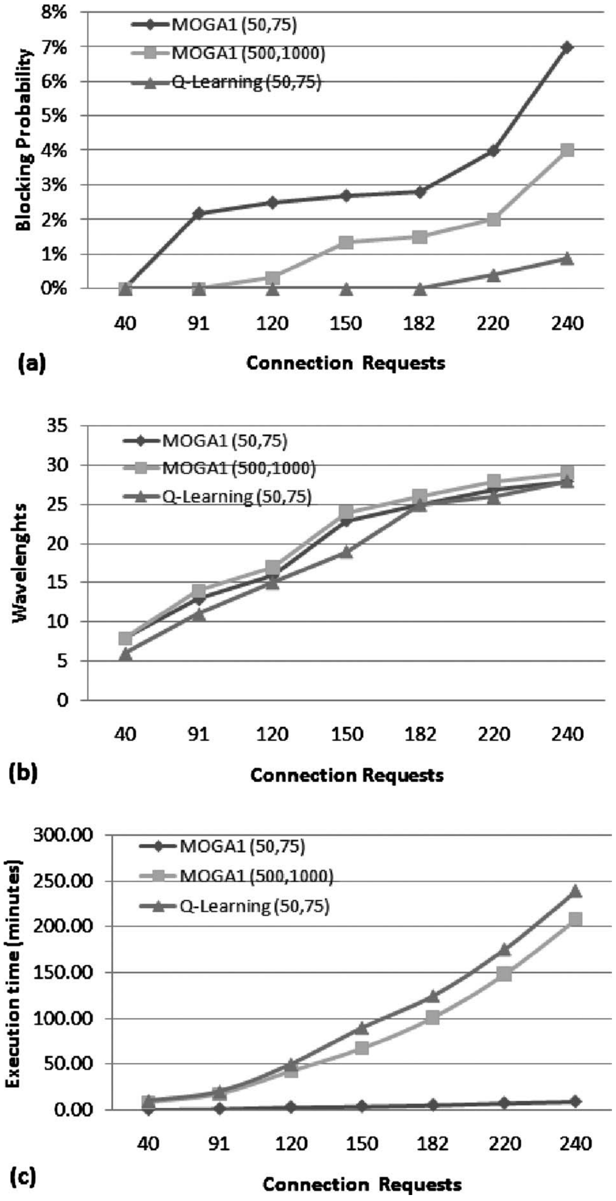

Fig. 7. Blocking ratio, number of wavelengths, and execution time versus connection requests for the MOGA and the Q-learning schemes.

7(a) shows the blocking performance, Fig. 7(b) shows the number of wavelengths required, and Fig. 7(c) shows the execution time.

From Fig. 7 it can be seen that the Q-learning scheme outperforms the MOGA scheme at the expense, however, of a significant increase in the execution time. In particular, zero blocking has been achieved for a fully connected network (one request for each source–destination pair), but the execution time is almost 2 hours and continues to increase with the increase of connection requests. This is because the Q-learning genetic algorithm estimates the Q factor at each iteration (epoch) in order to evaluate the fitness of the chromosome and evolve the population.

The case of the MOGA scheme with an initial/maximum population of 500/1000 is slightly better but again computationally heavy. The gain in blocking performance is 1%, and execution time is ~100 min. In contrast, the performance of the MOGA scheme with a 50/75 population size (also drawn in Fig. 7 for reference) is by far lighter, but, however, performs worse in terms of blocking. To this end, we may argue that the population size actually determines the MOGA performance. Hence, for a large number of individuals, the MOGA performance resembles that of the Q-based scheme, which can be considered to be the most accurate one due to the explicit Q-factor estimations. In contrast, having smaller populations result in less computational overhead but with a lower performance (higher blocking) as well. In any case, we may argue that performance and accuracy depends on the population size and can be decided on an application basis. For example, for the dynamic IA-RWA problem, one may tune the population size to meet traffic arrival requests and provision a certain blocking ratio.

V. CONCLUSION

In this paper, we presented two genetic algorithms for solving the IA-RWA algorithms. The first one is a heuristic, multiobjective genetic algorithm (MOGA) that accounts for fiber impairments indirectly through the insertion in the optimization process of the path length and number of common hops. The second algorithm is a Q-learning genetic algorithm that evaluates the Q factor at each iteration before evolving the population. The two schemes were evaluated and compared with each other. Results have shown that the MOGA scheme is capable of adequately handling fiber impairments, while its performance is determined by the population size. To this end, for a large number of individuals, performance is significantly improved and converges towards that of the Q-learning scheme, at the cost, however, of convergence speed. In contrast, by selecting small population sizes, speed is improved but results in a higher blocking ratio.

ACKNOWLEDGMENTS

The work described in this paper was carried out with the support of the BONE project ("Building the Future Optical Network in Europe"), a Network of Excellence funded by the European Commission through the 7th ICT-Framework Program.

REFERENCES

[1] B. Mukherjee, *Optical Communication Networks*. McGraw-Hill, 1997.

[2] S. Yan, M. Ali, and J. Deogun, "Route optimization of multicast sessions in sparse light-splitting optical networks," in *IEEE Globecom*, 2001, pp. 2134–2138.

[3] B. Van Caenegem, W. Van Parys, F. De Turck, and P. M. Demeester, "Dimensioning of survivable WDM networks," *IEEE J. Sel. Areas Commun.*, vol. 16, no. 7, pp. 1146–1157, 1998.

[4] N. Banerjee, V. Mehta, and S. Pandey, "A genetic algorithm approach for solving the routing and wavelength assignment problem in WDM networks," in *Proc. of Int. Conf. on Networking*, 2004.

[5] L. Zong and B. Ramamurthy, "Optimization of amplifier placement in switch-based optical network," in *Proc. of IEEE ICC*, 2001, vol. 1, pp. 224–228.

[6] B. Ramamurthy, D. Datta, H. Feng, J. P. Heritage, and B. Mukherjee, "Impact of transmission impairments on the teletraffic performance of wavelength routed optical networks," *J. Lightwave Technol.*, vol. 17, no. 10, pp. 1713–1723, 1999.

[7] T. Deng and S. Subramaniam, "Adaptive QoS routing in dynamic wavelength-routed optical networks," in *Proc. of Broadnets*, 2005, vol. 1, pp. 184–193.

[8] J. He, M. Brandt-Pearce, Y. Pointurier, and S. Subramaniam, "QoT-aware routing in impairment-constrained optical networks," in *Proc. of Globecom*, 2007, pp. 26–30.

[9] D. Monoyios, K. Vlachos, M. Aggelou, and I. Tomkos, "On the use of multi-objective optimization algorithms for solving the impairment aware-RWA problem," in *Proc. of ICC*, 2009.

[10] C. M. Fonseca and P. J. Fleming, "Multiobjective optimization and multiple constraint handling with evolutionary algorithms—part I: a unified formulation," *IEEE Trans. Syst. Man Cybern.*, vol. 28, no. 1, pp. 26–37, 1998.

[11] D. Brelaz, "New methods to color the vertices of a graph," *Commun. ACM*, vol. 22, pp. 251–256, 1979.

[12] http://www.vpisystems.com/.

[13] R. Kumar and P. Rockett, "Improved sampling of the Pareto-front in multiobjective genetic optimizations by steady-state evolution: a Pareto converging genetic algorithm," *Evol. Comput.*, vol. 10, no. 3, pp. 283–314, 2002.

Demetris Monoyios received the diploma in computer engineering from the Computer Engineering and Informatics Department (CEID) of the University of Patras, Greece, in 2007. He is currently working towards his Ph.D. degree. His research interests include optical networks, optical burst switching, algorithms, and protocols. Mr. Monoyios is the (co)author of three publications.

Kyriakos G. Vlachos received his Dipl.-Ing. degree in electrical and computer engineering from the National University of Athens (NTUA), Greece, in 1998 and his Ph.D. in electrical and computer engineering, also from NTUA, in 2001. From 1997 to 2001 he was a Senior Research Associate in the Photonics Communications Research Laboratory (ICCS/NTUA). In April 2001 he joined Bell Laboratories, Lucent Technologies, working on behalf of the Applied Photonics Group. Prof. Vlachos conducted research on high-speed optical networks and DWDM transmission techniques. During 2003, he joined the National Regulation Authority of Telecommunication and Postal Service of Greece (EETT), where he served as a Scientific Advisor for various techno-economic issues for the promotion of broadband technologies. Since 2003, he was also a member of the Computer Engineering Laboratory of the Technical University of Delft, and since 2005 he has been a Faculty Member of the Computer Engineering and Informatics Department of the University of Patras, Greece. In 2007, he founded the Photonic Networks and Technology Laboratory (PNET) (http://ru1photonicslab.cti.gr/). His research interests are in the areas of high-speed protocols and technologies for broadband, high-speed networks, optical packet/burst switching, and grid networks. Prof. Vlachos has participated in various research projects funded by the European Commission (IST-STOLAS, IST-PRO3, ESPRIT-DOALL, ephoton/ONe+, ICT-BONE, and ICT-DICONET). Prof. Vlachos is a member of IEEE and the Technical Chamber of Greece, and he periodically acts as a scientific reviewer for the General Secretariat for Research and Technology of Greece (GSRT) as well as for the European Commission and the Netherlander Organization for Scientific Research, Technology Foundation. Prof. Vlachos is the (co)author of more than 90 journal and conference publications and holds 5 patents.