

Fair Resource Allocation in Optical Networks under Tidal Traffic

Tania Panayiotou and Georgios Ellinas

KIOS Research and Innovation Center of Excellence

Department of Electrical and Computer Engineering, University of Cyprus

{panayiotou.tania, gellinas}@ucy.ac.cy

Abstract—We propose an α -fair routing and spectrum allocation (RSA) framework for reconfigurable elastic optical networks under modeled tidal traffic, that is based on the maximization of the social welfare function parameterized by a scalar α (the inequality aversion parameter). The objective is to approximate an egalitarian spectrum allocation (SA) that maximizes the minimum possible SA over all connections contending for the network resources, shifting from the widely used utilitarian SA that merely maximizes the network efficiency. A set of existing metrics are examined (i.e., connection blocking, resource utilization, coefficient of variation (CV) of utilities), and a set of new measures are also introduced (i.e., improvement on connection over- (COP) and under-provisioning (CUP), CV of unserved traffic), allowing a network operator to derive and evaluate in advance a set of α -fair RSA solutions and select the one that best fits the performance requirements of both the individual connections and the overall network. We show that an egalitarian SA better utilizes the network resources by significantly improving both COP (up to 20%) and CUP (up to 80%), compared to the utilitarian allocation, while attaining zero blocking. Importantly, the CVs of utilities and unserved traffic indicate that a SA that is fairest with respect to the amount of utilities allocated to the connections does not imply that the SA is also fairest with respect to the achievable QoS of the connections, while an egalitarian SA better approximates a fairest QoS-based SA.

Index Terms—fairness, game theory analysis, optimization, routing and spectrum allocation, tidal traffic, optical networks

I. INTRODUCTION

Allocation of optical network resources has long been associated with the objective of maximizing the *network efficiency*, with the efficiency usually measured according to performance evaluation metrics reflecting the aggregated network throughput (e.g., connection blocking, spectrum utilization, etc.). Maximization of the network efficiency, however, lacks *fairness*, as it does not ensure that the available resources are fairly divided among the connections to achieve a similar quality-of-service (QoS) for all connections. Specifically, in a network environment, where connections are likely to contend for the same bandwidth resources, maximization of the network efficiency tends to favor connections with higher bit-rate tendencies, possibly blocking connections with lower bit-rate tendencies. Furthermore, the commonly adopted performance evaluation metrics do not account for the examination of fairness measurements, and hence optimization of such measurements is not taken into consideration during resource allocation.

On this basis, in this work we examine fairness considerations according to an α -fair resource allocation scheme [1] that allows a number of fairness measures to be explicitly examined for a reconfigurable elastic optical network (EON) under tidal traffic. In general, tidal traffic refers to the repetitive spatio-temporal variations of traffic demand (i.e., fluctuations), caused mainly by the daily large-scale population migration among different areas at different times [2], and enhanced by the emergence of new types of networks (e.g., 5G networks), applications (e.g., video streaming), and services (e.g., connected cars). As tidal traffic is present in optical transport segments, SDN-enabled reconfigurable optical networks capable of following as closely as possible tidal traffic, have recently gained attention from the research community [3]–[12].

Specifically, related work deals with traffic demand prediction [3]–[7], [9], [10], [12] and with predictions used as inputs to resource allocation problems [3]–[12], allowing in advance network optimization. The performance evaluation metrics commonly considered are associated with the network throughput, that is measured in connection blocking percentage [3], [9]–[12], spectrum utilization (e.g., number of frequency slots utilized) [9]–[12], and/or the degree of defragmentation [11]. However, such performance evaluation metrics lack fairness information, especially as it concerns the QoS of contending connections. Moreover, the commonly considered optimization objectives do not account for the fairness of the allocations. Only a few works exist in the literature focusing on fair resource allocation strategies [13]–[15]. However, in these works, the only measure of fairness considered is the blocking probability.

To this end, we propose the development of an α -fair routing and spectrum allocation (RSA) framework for reconfigurable EONs, under modeled tidal traffic, aiming to derive and evaluate in advance fair SAs. In this framework, an SDN-based central decision management unit (i.e., network orchestrator) decides on the allocation by maximizing the constant elasticity social welfare function $W_\alpha(u)$, parameterized by $\alpha \geq 0$ and defined for $u \in \mathbb{R}_+^n$, where u is an α -fair utility allocation for the n players. In this work, the connections are perceived as the players competing for the link utilities, and the utilities are the frequency slots (FSs) available on the network links. The α -fair RSA is formulated as an Integer Linear Program (ILP) that can be executed in advance, for each (re)configuration

time interval of interest, allowing the reconfiguration to take place at the beginning of each time interval.

Even though α is considered to be a natural measure of fairness [1], in this work the α -fair RSA is additionally evaluated according to several other, easier to interpret, measures, that also facilitate an understanding of the efficiency-fairness trade-offs for different α values. These measures include connection blocking percentage, resource utilization, and the CV of the utilities [16], as well as other measures that are introduced in this work aiming to directly interpret the fairness of the spectrum allocations (SAs) with respect to the QoS of contending connections. In particular, we introduce the improvement on connection over- (ICOP) and under-provisioning (ICUP), and the CV of the expected unserved traffic. These measures can be evaluated in advance, for several α values, according to the fluctuations resulting from the modeled tidal traffic. This allows the centralized controller to decide the α -fair RSA solution that best meets the performance requirements of both the connections and the network, before network (re)configuration takes place.

Overall, we show that as α increases, the proposed ILP-based α -fair RSA algorithm approximates an egalitarian SA (i.e., an allocation that maximizes the minimum possible SA over all connections contending for the network resources) by significantly improving connection blocking, COP and CUP, compared to the widely used utilitarian allocation (i.e., an allocation that maximizes the network efficiency). Furthermore, the CV of utilities indicates that a SA that is fairest with respect to the amount of allocated utilities, does not imply that the SA is also fairest with respect to the QoS that contending connections enjoy. Importantly, the CV of unserved traffic indicates that an egalitarian allocation best approximates a SA that is fairest with respect to the connections' QoS. It is worth mentioning that α -fair schemes have been previously examined in wireless networks [17], SDN IP networks [18], and TCP networks [19]–[21]. In optical networks, however, α -fairness has not been previously considered.

II. α -FAIRNESS PRELIMINARIES

According to the α -fair scheme, a central decision management unit decides on the allocation of the available resources among n players by maximizing the constant elasticity welfare function W_α parameterized by $\alpha \geq 0$ [1] as

$$W_\alpha(u) = \begin{cases} \sum_{i=1}^n \frac{u_i^{1-\alpha}}{1-\alpha} & \text{for } \alpha \geq 0, \alpha \neq 1 \\ \sum_{i=1}^n \log(u_i) & \text{for } \alpha = 1 \end{cases} \quad (1)$$

where $u \in \mathbb{R}_+^n$ is a *utility allocation* and u_i is equal to the utility derived by player i . In general, a utility allocation is *feasible* if and only if there exists an allocation of resources for which the utilities derived by the players are u_1, u_2, \dots, u_n , accordingly. The utility set $\mathcal{U} \subset \mathbb{R}_+^n$ is then defined as the set of all feasible utility allocations, and a resulting α -fair allocation, denoted by $\pi(\alpha)$, is such that $\pi(\alpha) \in \arg\max_{u \in \mathcal{U}} W_\alpha(u)$.

An explanation of why W_α yields fair allocations is that it exhibits diminishing marginal welfare increase, as utilities increase (because W_α is concave and component-wise increasing), with parameter α controlling the rate at which marginal increases diminish [1]. For example, if player i is allocated less utilities than player j , then a marginal increase in the utilities of player i (as α increases) would yield a higher welfare increase compared to a marginal increase in the utilities of player j . Hence, the marginal increase in the utilities of player i would be more preferable.

Special cases of the α -fairness scheme are derived for $\alpha = 0$, $\alpha = 1$, and $\alpha \rightarrow \infty$. Specifically, for $\alpha = 0$ the W_α function corresponds to finding a *utilitarian allocation* of the associated optimization problem; that is, an optimal solution of the problem that results in the highest system efficiency (e.g., maximum network throughput). For $\alpha = 1$ the scheme corresponds to *proportional fairness* (the generalization of the Nash solution for a two-player problem) [22], and for $\alpha \rightarrow \infty$ the scheme converges to *max-min fairness*; that is, an egalitarian allocation that ensures that the minimum allocated utility is as high as possible. Hence, a higher α value corresponds to a fairer scheme [1]. It is worth noting that utilitarian, proportional, and egalitarian schemes yield Pareto optimal allocations [22]. For our network application scenario, Pareto optimality means that the connections that do not contend with other connections along their routes (links) will be allocated their maximum requested rates, subject to the link capacity constraints.

III. PROBLEM STATEMENT

Given an EON with k links and M FSs at each network link, the general objective is to find an α -fair RSA for n connection requests. For solving the conventional RSA problem, a feasible route and a SA must be found for each connection request, forming a lightpath. For the proposed α -fair RSA the assumption is that the routes are pre-computed and given in the link utilization matrix $P = [p_{il}] \in \mathbb{R}^{n \times k}$, indicating the contending connections for each link in the network (i.e., the connections utilizing the same links). Therefore, in this work, the utilities are the available FSs, and the players are the connections. The utilities are given in the SA matrix $U = [u_{ij}] \in \mathbb{R}_+^{n \times m}$, consisting of all the possible SAs (actions) for each connection i . Hence, u_{ij} is SA j for connection i , measured in the number of FSs, and m is the number of possible SAs for each connection i . Since M is the available link bandwidth, then $0 \leq u_{ij} \leq M$, for all connections and SAs.

Given matrices P and U , the *objective* is to obtain an egalitarian α -fair SA subject to the feasibility constraints of the EON technology. The problem is formulated as an ILP with the objective of maximizing the social elasticity welfare function, $W_\alpha(\hat{u})$, parameterized by $\alpha \geq 0$ and defined for $\hat{u} \in \mathbb{R}_+^n$, where $\hat{u} = [\hat{u}_1, \dots, \hat{u}_n] \in \hat{\mathcal{U}}$ is a feasible SA of the normalized utilities $\hat{u}_i \in \hat{\mathcal{U}}$, derived by normalizing the SA matrix U . Note that, by normalizing the utilities we capture the fact that we are less concerned with the amount of utility achieved by a connection and more with the fraction

of maximum possible utility that a connection achieves (i.e., the satisfaction of an end-user); a common practice in welfare maximization problems that focus on approximating egalitarian welfare [23]. Hence, an α -fair allocation of the normalized utilities is given by $\hat{\pi}(\alpha) \in \arg\max_{\hat{u} \in \hat{\mathcal{U}}} W_\alpha(\hat{u})$, from which the α -fair allocation of the utilities, $\pi(\alpha)$, can be derived. Note that sets \mathcal{U} and $\hat{\mathcal{U}}$ are derived during the execution of the ILP algorithm.

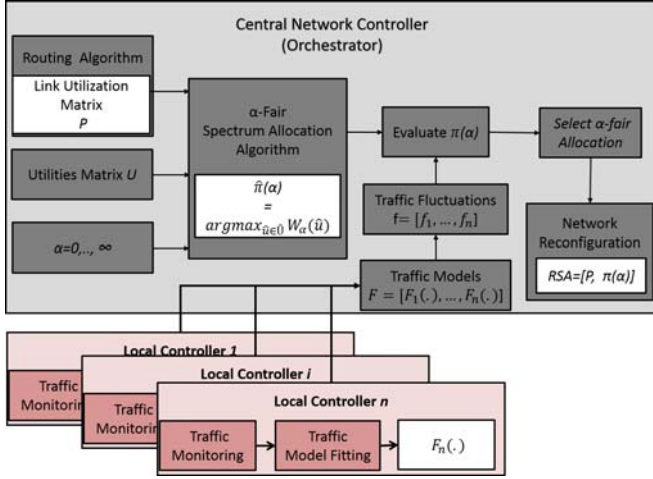


Fig. 1. α -fair RSA framework.

Figure 1 illustrates the proposed α -fair RSA framework. This framework is performed off-line, for each time interval of interest (e.g., hourly, daily time intervals), and the network configurations are readily available at the beginning of each time interval. For simplicity, time interval notations are omitted from Fig. 1. The α -fair RSA algorithm is executed for several α values, and each α -fair allocation $\pi(\alpha)$ is evaluated in advance. Hence, the network operator is capable of selecting the α -fair RSA that best meets specific performance targets.

Specifically, α -fair allocations are evaluated according to the connection blocking percentage, resource utilization, CV of utilities, CV of expected unserved traffic, and the improvement on COP/CUP. The latter measures are based on the expected unserved traffic and excess utilities, $u^- = [u_1^-, \dots, u_n^-]$ and $u^+ = [u_1^+, \dots, u_n^+]$, respectively, estimated according to the a-priori modeled tidal traffic. Tidal traffic, is given in a set of traffic demand distributions $F = [F_1(\cdot), \dots, F_n(\cdot)]$, from which traffic fluctuations $f = [f_1, \dots, f_n]$ are drawn as $f \sim F$. Hence, u^- and u^+ are estimated by the fluctuations around $u \in \pi(\alpha)$. Note that each distribution is assumed to be a function of the bandwidth demand measured according to the requested FSs. Even though the requested number of FSs depends on several factors, such as the modulation format, in this work, similar to [10], for simplicity we assume that the distributions directly reflect the requested number of FSs (i.e., logical link usage); an assumption that does not affect the scope of this work.

According to Fig. 1, traffic demand models, F , are present in local controllers (i.e., offloading the central controller from

functionalities that can be performed locally) by means of monitoring and analyzing the aggregated traffic of each source-destination pair forming a logical connection. After model fitting, the relevant parameters (e.g., distribution, mean, and standard deviation) are communicated to the central controller, from which fluctuations f can be drawn. The reader should note here that the traffic modeling procedure is omitted as it is out of the scope of this work.

IV. α -FAIR ROUTING AND SPECTRUM ALLOCATION

The α -fair RSA algorithm is based on the maximization of $W_\alpha(\hat{u})$, where $\hat{u} = [\hat{u}_1, \dots, \hat{u}_n] \in \hat{\mathcal{U}}$, and $\hat{u}_i \in \hat{\mathcal{U}}$ are the normalized utilities of SA matrix U . Specifically,

• $U = [u_{ij}] \in \mathbb{R}^{n \times m}$ is the matrix describing all the possible SA j of each connection i , for a total of m possible SAs, and

$$u_{ij} = \begin{cases} j \times \frac{M}{m} & \text{if } (j \times \frac{M}{m}) \leq u_{imax} \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

where u_{imax} is the peak-rate demand of connection i (i.e., given by $F_i(\cdot)$). According to Eq. (2), the number of FSs increases as index j increases, with a maximum possible allocation equal to the maximum link capacity, M . For each connection i , u_{ij} is set to 0 when $j \times M/m$ is greater than its peak-rate demand. By doing so, we avoid including in U allocations that will lead, with certainty, to unnecessary connection overprovisioning.

• $\hat{U} = [\hat{u}_{ij}] \in \mathbb{R}^{n \times m}$ is the SA matrix of the normalized utilities, where

$$\hat{u}_{ij} = \begin{cases} \frac{u_{ij}}{u_{imax}} & \text{if } u_{ij} > 0 \\ \epsilon & \text{otherwise} \end{cases} \quad (3)$$

and ϵ is a small positive value used for ensuring that $\hat{u} \in \mathbb{R}_+^n$. Note that ϵ must be greater than 0 and less than the smallest possible SA (i.e., $\epsilon \ll M/m$). The denominator u_{imax} is used for normalizing the utilities, such that W_α takes into account the normalized utilities $\hat{u} \in [\epsilon, 1]$.

For the ILP formulation, we assume that the routing problem is solved using Dijkstra's algorithm [24], and

• $P = [p_{il}] \in \{0, 1\}^{n \times k}$, $i = 1, \dots, n$, $l = 1, \dots, k$ is the link utilization matrix formed according to the pre-computed routes; $p_{il} = 1$ if connection i utilizes link l , and 0 otherwise. The variables, objective, and constraints of the ILP are defined as follows:

Variables:

- x_{ij} : Boolean variable equal to 1 if SA j is chosen for connection i , and 0 otherwise.
- y_{is} : Boolean variable equal to 1 if FS s is utilized by connection i , and 0 otherwise.
- z_{is} : Boolean variable equal to 1 if FS s is the first FS utilized by connection i amongst a set of contiguous FSs allocated to connection i , and 0 otherwise.

Objective:

$$\text{Maximize: } W_\alpha(\hat{u}) = \begin{cases} \sum_{ij} x_{ij} \frac{\hat{u}_{ij}^{1-\alpha}}{1-\alpha} & \text{if } \alpha \geq 0, \alpha \neq 1 \\ \sum_{ij} x_{ij} \log(\hat{u}_{ij}) & \text{if } \alpha = 1 \end{cases} \quad (4)$$

Subject to:

$$\sum_{ij} u_{ij} x_{ij} p_{il} \leq M, \quad \forall l = 1, \dots, k \quad (5)$$

$$\sum_j x_{ij} \leq 1, \quad \forall i = 1, \dots, n \quad (6)$$

$$\sum_s y_{is} = \sum_j u_{ij} x_{ij}, \quad \forall i = 1, \dots, n \quad (7)$$

$$\sum_i y_{is} p_{il} \leq 1, \quad \forall s = 1, \dots, M, l = 1, \dots, k \quad (8)$$

$$\sum_s z_{is} \leq 1, \quad \forall i = 1, \dots, n \quad (9)$$

$$y_{is} - y_{i(s-1)} \leq z_{is}, \quad \forall i = 1, \dots, n, s = 2, \dots, M \quad (10)$$

$$y_{is} = 0 \quad \text{for } s = 1, \quad \forall i = 1, \dots, n \quad (11)$$

The objective of the ILP is to maximize $W_\alpha(\hat{u})$, resulting in an α -fair allocation $\hat{\pi}(\alpha) = \arg\max_{\hat{u} \in \hat{\mathcal{U}}} W_\alpha(\hat{u})$ of the normalized utilities, from which the actual α -fair SA $\pi(\alpha) = u \in \mathcal{U}$ can be derived. Constraint (5) ensures that the SA decisions meet the link capacity constraints, while constraint (6) ensures that at most one allocation is chosen for each connection. Constraint (7) ensures that the number of allocated FSs is equal to the number of FSs indicated by the SA implemented and constraint (8) ensures that the no-frequency overlapping constraint is met for all network links. Finally, constraints (9), (10), and (11) ensure that each connection is assigned contiguous FSs on all links of its route.

V. EVALUATION MEASURES

Once the α -fair RSA is derived, the expected unserved traffic and excess utilities $u^- = [u_1^-, \dots, u_n^-]$ and $u^+ = [u_1^+, \dots, u_n^+]$, respectively, are computed as

$$u_i^+ = \frac{1}{T} \sum_{t|u_i > f_{it}} (u_i - f_{it}) \quad (12)$$

$$u_i^- = \frac{1}{T} \sum_{t|u_i < f_{it}} |(u_i - f_{it})| \quad (13)$$

where $u_i \in \pi(\alpha)$ and $f_{it} \sim F_i(\cdot)$, $\forall t = 1, \dots, T$. Hence, u^+ and u^- are evaluated according to T traffic demand fluctuation around the α -fair utilities $u \in \pi(\alpha)$.

COP and CUP are introduced as measures of the aggregated expected excess utilities and unserved traffic, respectively, and are given by $COP(\alpha) = \sum_i u_i^+$ and $CUP(\alpha) = \sum_i u_i^-$. Improvements on COP (ICOP) and CUP (ICUP) are given by

$$ICOP(\alpha) = \frac{COP(0) - COP(\alpha)}{COP(0)}, \quad ICUP(\alpha) = \frac{CUP(0) - CUP(\alpha)}{CUP(0)} \quad (14)$$

where $COP(0)$ and $CUP(0)$ are derived by the utilitarian allocation $\pi(0)$, which is likely to cause the highest values of COP and CUP. As α increases (a fairer allocation is derived), $ICOP(\alpha)$ is likely to only positively increase (up to one), as the contesting connections with higher bit-rate tendencies (over-provisioned) are likely to be allocated less resources, hence improving the overall COP. On the other hand, as α increases, $ICUP(\alpha)$ is likely to positively increase (as more connections are admitted into the network), without any guarantees however that it will not start to decrease as $\alpha \rightarrow \infty$ (e.g., max-min fairness is achieved by increasing under-provisioning). The latter depends on the network load and the traffic demand behavior of each connection. Nevertheless, the evaluation of these measures allows the network operator to decide on the allocation that best meets specific performance criteria for the connections and the overall network.

Furthermore, even though parameter α is considered to be a natural measure of fairness, in this work we also use the CV of the utilities [16], to directly interpret the fairness of the allocations with respect to the amount of utilities (FSs) allocated to the connections. In particular, the CV of the utilities measures the dispersion of utilities allocated to the connections around the mean of the allocated utilities and is given by

$$CV(\alpha) = \sqrt{\frac{1}{n-1} \sum_{i=1}^n \frac{(u_i - \bar{u})^2}{\bar{u}^2}} \quad (15)$$

where $u_i \in \pi(\alpha)$, and $\bar{u} = \frac{1}{n} \sum_{i=1}^n \pi_i(\alpha)$ is the mean of allocated utilities. An α -fair allocation is said to be fairer with respect to the utilities, if and only if $CV(\alpha)$ is smaller than the CVs of other α values. Hence, as α increases, $CV(\alpha)$ is likely to decrease. As we are also interested in investigating the fairness of the allocations with respect to the connections' achievable QoS, we also use the CV of the unserved traffic, evaluated according to Eq. (15) as a function of u^- , instead of u . The α -fair allocation that minimizes the CV of u^- is the one that best approximates the fairest allocation with respect to the connections' QoS. Note that it is more critical to derive a fair u^- , instead of a fair u^+ , as connection under-provisioning is a noticeable and annoying effect for the end-users, while connection over-provisioning just tends to waste resources. Finally, regarding connection blocking, a connection is considered blocked when allocated 0 FSs (i.e., $u_i = 0$), and resource utilization is measured as the sum of the FSs (utilities) allocated to all connections along their routes.

VI. PERFORMANCE EVALUATION

In order to evaluate the performance of the proposed framework, the generic Deutsche Telekom (DT) network (14 nodes, 23 links) of Fig. 2 was used, with each FS set at 12.5GHz and with each fiber link utilizing $M = 100$ FSs. In total $m = 50$ possible SAs were assumed for each one of the $n = 20$ connections, with randomly selected source-destination pairs.

For the evaluation of u^+ and u^- , $T = 1000$ fluctuations (f) were drawn from a set of log-normal distributions F .




Fig. 2. Deutsche Telekom network topology.

Specifically, for each connection i , $f_i \sim F(\mu_i, \sigma_i^2)$, with each $f_i \in (0, M)$. Parameters σ_i^2 and μ_i , were drawn from uniform distributions $unif(0, 1)$ and $unif(2.5, 4.5)$, respectively. For the simulations, fluctuations f_i were scaled down by a factor of two in order to fit the link capacity constraints. Further, the scaled fluctuations indicating a higher spectrum demand than M were set to $f_i = M$. It is important to note that the log-normal distribution is used, as it has been shown that Internet traffic volumes are best characterized by such a distribution [25], [26] (i.e., by skewness and heavy tails).

The α -fair RSA ILP is solved in a MATLAB machine with a CPU @2.60GHz and 8GB RAM. Parameter α was examined for the range of values $[0, 5]$ with a 0.1 step, resulting in 50 α -fair allocations $\pi(\alpha)$. In total, 7.5 hours were required to solve all 50 ILPs, with the running time of each ranging between 1.5 hrs and 50 sec. The resulting α -fair allocations were evaluated according to the connection blocking percentage and resource utilization (Fig. 3), COP and CUP (Fig. 4), ICOP and ICUP (Fig. 5), and CVs of u and u^- (Fig. 6).




Fig. 3. (a) Connection blocking percentage vs α . (b) Resource utilization in number of FSs vs α .




Fig. 4. (a) $COP(\alpha)$ in number of FSs. (b) $CUP(\alpha)$ in number of FSs.

The results indicate that as α increases, connection blocking percentage reduces (Fig. 3(a)). The highest connection




Fig. 5. (a) $ICOP(\alpha)$. (b) $ICUP(\alpha)$.




Fig. 6. (a) $CV(\alpha)$ of spectrum allocations u . (b) $CV(\alpha)$ of unserved traffic u^- .

blocking is observed for the utilitarian allocation ($\alpha = 0$). This is reasonable, as for $\alpha = 0$ the objective is to maximize the aggregated network throughput, favoring the connections with higher bit-rate tendencies, and consequently blocking connections with lower bit-rate tendencies. Even though zero connection blocking is achieved for $\alpha = 0.3$, results regarding resource utilization, COP/CUP, and CV of u and u^- , suggest that α values greater than 0.3 achieve fairer allocations; an indicator that zero connection blocking alone does not imply fairness.

In particular, Fig. 3(b) suggests that as α increases, the available spectrum is more efficiently utilized, as the resources required for serving all the connections keep decreasing. Furthermore, as α increases, both COP and CUP tend to decrease (Fig. 4), striking a better balance between the COP and CUP effects. Specifically, for the larger α values examined, COP improves by approximately 20% (Fig. 5(a)) and CUP improves by approximately 80% (Fig. 5(b)). In particular, for our problem setting, approximately the highest improvement in both ICUP and ICOP is achieved around $\alpha = 1$, indicating that proportional fairness can be considered as a sufficient solution (if one considers only the ICUP/ICOP measures).

Furthermore, Fig. 6 shows that the CVs of both u and u^- tend to decrease as α increases, indicating that the SAs are fairer with respect to both the utilities u and the unserved traffic u^- . It is important to note that the CVs of both unserved traffic and utilities do not change for α values greater than 4.5 (although not shown in the results α values greater than 5 were investigated), as an egalitarian SA is reached (i.e., the minimum allocated utility did not further increase). Note that the slight increase in the CV of utilities (Fig. 6(a)) for α values greater than 1, occurs due to the non-contending connections

that are always allocated their peak-rate requests, consequently causing a variation of utilities that is slightly larger around the mean allocated utilities, as the variability of the utilities of contending connections decreases. Further, the increase in the CV of unserved traffic (Fig. 6(b)) for α values ranging from 0 to 0.2, is caused due to the blocked connections in that range (Fig. 3) that differently affect the CV of unserved traffic depending on the traffic demand behavior (i.e., mean and standard deviation) of the exact connections that were blocked.

Importantly, the CVs of both the utilities and the unserved traffic allow the network operator to decide which α -fair allocation best fits specific performance requirements. In particular, Fig. 6(b) suggests that α values greater than 2 best approximate fairer allocations with respect to u^- , while Fig. 6(a) suggests that when $\alpha = 1$ (i.e., proportional fairness), an allocation that is fairer with respect to u is achieved. Note that proportional fairness also achieves the highest ICUP (Fig. 5(b)). As, however, the network operator may be more interested in approximating an egalitarian SA, achieving at the same time, as similar as possible levels of u_i^- , α values greater than 2 may be preferable. Specifically, $\alpha \geq 2$ achieves zero connection blocking (Fig. 3(a)), only affects resource utilization by approximately 5% (Fig. 3(b)), and achieves approximately 20% ICOP and 80% ICUP (Fig. 5). Furthermore, CV of u indicates that the fairest SA with respect to the amount of allocated utilities does not imply that the SA is fairest with respect to the connections' QoS. CV of u^- , on the other hand, indicates that the fairest QoS-based SA is best approximated by an egalitarian SA (i.e., CV of u^- decreases as $\alpha \rightarrow \infty$).

VII. CONCLUSION

In this work, we proposed an α -fair RSA framework that is based on the maximization of the social welfare function of the normalized utilities and allows a set of α -fair allocations to be computed and evaluated in advance for reconfigurable EONs under modeled tidal traffic. This framework allows for the examination of several existing measures (i.e., connection blocking, resource utilization, CV of utilities), as well as the emergence of new measures (i.e., CV of unserved traffic, COP/CUP, ICOP/ICUP) and their exploitation by a network operator so as to decide on the α -fair allocation that best fits the performance criteria of both the connections and the network. Importantly, we show that the new measures and the α -fair RSA algorithm achieve to approximate an egalitarian SA that it is also fairer with respect to the QoS for all connections. The egalitarian SA significantly improves connection blocking, expected COP (up to 20%), and CUP (up to 80%), as opposed to the widely used utilitarian allocation. The development of α -fair heuristic approaches for larger problem sizes, as well as techniques for jointly solving the α -fair RSA problem while also considering of the various modulation formats are planned as future work.

ACKNOWLEDGMENT

This work has been supported by the European Union's Horizon 2020 research and innovation programme under grant agreement No 739551 (KIOS CoE) and from the Government of the Republic of Cyprus through the Directorate General for European Programmes, Coordination and Development.

REFERENCES

- [1] D. Bertsimas, et al., "On the efficiency-fairness trade-off", *Management Science*, 58(12):2234–2250, 2012.
- [2] Z. Zhong et al., "Energy efficiency and blocking reduction for tidal traffic via stateful grooming in IP-over-optical networks", *IEEE/OSA J. Opt. Commun. Netw.*, 8(3):175–189, 2016.
- [3] B. Yan, et al., "Tidal-traffic-aware routing and spectrum allocation in elastic optical networks", *IEEE/OSA J. Opt. Commun. Netw.*, 10(11):832–842, 2018.
- [4] R. Alvizu, et al., "Matheuristic with machine-learning-based prediction for software-defined mobile metro-core networks", *IEEE/OSA J. Opt. Commun. Netw.*, 9(9):D19–D30, 2017.
- [5] F. Morales, et al., "Virtual network topology adaptability based on data analytics for traffic prediction", *IEEE/OSA J. Opt. Commun. Netw.*, 9(1):A35–A45, 2017.
- [6] N. Fernández et al., "Virtual topology reconfiguration in optical networks by means of cognition: Evaluation and experimental validation", *IEEE/OSA J. Opt. Commun. Netw.*, 7(1):A162–A173, 2015.
- [7] Y. Ohsita et al., "Gradually reconfiguring virtual network topologies based on estimated traffic matrices", *Proc. IEEE INFOCOM*, 2007.
- [8] A. Caballero et al., "Cognitive, heterogeneous and reconfigurable optical networks: The CHRON project", *IEEE/OSA J. Light. Techn.*, 32(13):2308–2323, 2014.
- [9] T. Panayiotou, et al., "On learning bandwidth allocation models for time-varying traffic in flexible optical networks", *Proc. ONDM*, 2018.
- [10] T. Panayiotou, et al., "A data-driven bandwidth allocation framework with QoS considerations for EONs", *IEEE/OSA J. Light. Techn.*, 37(9):1853–1864, 2019.
- [11] P. Kokkinos, et al., "Pattern-driven resource allocation in optical networks", *IEEE Trans. Netw. Serv. Mgt.*, 16(2):489–504, 2019.
- [12] S.K. Singh and A. Jukan, "Machine-learning-based prediction for resource (re)allocation in optical data center networks", *IEEE/OSA J. Opt. Commun. Netw.*, 10(10):D12–D28, 2018.
- [13] F. Callegati, et al., "Service fairness in flexible optical networks" *Proc. OFC/NFOEC*, 2013.
- [14] S. Ma, et al., "A fairness-aware dynamic spectrum allocation scheme in elastic optical networks," *Proc. OECC*, 2013.
- [15] H. Chen, et al., "A joint fairness-aware and fragmentation-reduction spectrum allocation scheme in elastic optical networks," *Proc. OECC and PGC*, 2017.
- [16] R.K. Jain, et al., "A quantitative measure of fairness and discrimination", *DEC Technical Report DEC-TR-301*, pp. 1–38, 1984.
- [17] E. Altman, et al., "Generalized α -fair resource allocation in wireless networks", *Proc. IEEE CDC*, 2008.
- [18] Z. Allybokus, et al., "Multi-path alpha-fair resource allocation at scale in distributed software-defined networks", *IEEE J. Sel. Areas in Commun.*, 36(12):2655–2666, 2018.
- [19] A. Tang, et al., "Counter-intuitive throughput behaviors in networks under end-to-end control", *IEEE Trans. on Netw.*, 14(2):355–368, 2006.
- [20] T. Lan, et al., "An axiomatic theory of fairness in network resource allocation", *Proc. IEEE INFOCOM*, 2010.
- [21] L. Buzna, et al., "Controlling congestion on complex networks: Fairness, efficiency and network structure", *Sci. Rep.*, 7(9152), 2017.
- [22] D. Bertsimas, et al., "The price of fairness," *Oper. Res.*, 59(1):17–31, 2011.
- [23] H. Aziz, "Justifications of welfare guarantees under normalized utilities," *arXiv:1905.10774*, 2019.
- [24] T. H. Cormen, et al., *Introduction to algorithms* (Third ed.), "Section 24.3: Dijkstra's algorithm", MIT Press, 2009.
- [25] I. Antoniou, et al., "On the log-normal distribution of network traffic", *Physica D: Nonlinear Phenomena*, 167(1–2):72–85, 2002.
- [26] M. Kassim, et al., "Statistical analysis and modeling of Internet traffic IP-based network for tele-traffic engineering", *ARN J. of Eng. and Applied Sciences*, 10(3):1505–1512, 2015.