

To be Fair or Efficient or a Bit of Both^{*}

Moshe Zukerman^{1,*}, Musa Mammadov², Liansheng Tan³

Iradj Ouveysi¹ and Lachlan Andrew⁴

Abstract

Introducing a new concept of (α, β) -fairness, which allows for a bounded fairness compromise, so that a source is allocated a rate neither less than $0 \leq \alpha \leq 1$, nor more than $\beta \geq 1$, times its fair share, this paper provides a framework to optimize efficiency (utilization, throughput or revenue) subject to fairness constraints in a general telecommunications network for an arbitrary fairness criterion and cost functions. We formulate a non-linear program (NLP) that finds the optimal bandwidth allocation by maximizing efficiency subject to (α, β) -fairness constraints. This leads to what we call an efficiency-fairness function, which shows the benefit in efficiency as a function of the extent to which fairness is compromised. To solve the NLP we use two algorithms. The first is a well known branch-and-bound-based algorithm called Lipschitz Global Optimization and the second is a recently developed algorithm called Algorithm for Global Optimization Problems (AGOP).

We demonstrate the applicability of the framework to a range of example from sharing a single link to efficiency fairness issues associated with serving customers in remote communities.

Key words: non-linear programming, utility optimization, fairness, efficiency-fairness tradeoff, bandwidth allocation.

1 Introduction

Efficiency-fairness tradeoffs have been of interest to people from many walks of life [3–5,11,12,14,21,25,37,38]. Different societies and countries make their choices on these tradeoffs. This paper focuses on such tradeoffs in the context of telecommunications networks and provides a framework for evaluation and presentation of such tradeoffs. It applies to any network, topology and any fairness criterion. It can be applied to networks such as Resilient Packet Rings (RPR) [8,13] (the IEEE 802.17 standard for metropolitan area networks), local area networks (wireline and wireless) and wide area networks and decisions on what resources to provide to customers being cross-subsidised, such as those in remote communities. Fairness criteria compatible with the framework include *RIAS fairness* [9,10], *max-min fairness* [2,6,32], *proportional fairness* [15,16,25], *general weighted (GW) fairness* [35,36], and *minimum potential delay fairness* [18].

We introduce the concept of (α, β) -*fairness*. In particular, we define capacity assignment to be (α, β) -*fair* if the rate allocated to a flow is neither less than α times its fair allocation for $0 \leq \alpha \leq 1$, nor it is higher than β times its fair allocation for $1 \leq \beta$. In this paper we provide a framework for maximizing the *efficiency*, under a

* This research has been partially supported by National Natural Science Foundation of China under Grant No. 60473085 and partially by the Australian Research Council.

* Corresponding author.

¹ M. Zukerman and I. Ouveysi are with the Australian Research Council Special Research Centre for Ultra-Broadband Information Networks (CUBIN) an affiliated program of National ICT Australia, Electrical and Electronic Engineering (EEE) Department, The University of Melbourne, Victoria 3010, Australia (e-mail: mzu@unimelb.edu.au, iradjouveysi@yahoo.co.uk).

² M. Mammadov is with the School of Information Technology and Mathematical Sciences, University of Ballarat, Victoria 3353, Australia (email: m.mammadov@ballarat.edu.au).

³ L. Tan is with the Computer Science Department of Central China Normal University, Wuhan 430079, PR China (e-mail: l.tan@mail.ccnu.edu.cn).

⁴ L. Andrew is with the Department of Computer Science, Caltech, Pasadena, CA 91125, USA (email: lachlan@caltech.edu).

constraint of, say, (90%, 150%)-fairness. By efficiency we mean the value of a general utility function of the flow rate allocation. Examples of such utility functions are profit, revenue, utilization and throughput. In contrast to the approach used in [26] and [34], we do not seek to choose a particular utility function as a trade off between efficiency and fairness. In our framework, the utility function that defines efficiency is unrelated to the fairness criterion. Our framework addresses the question: How much can the efficiency be improved by compromising on fairness to a certain extent? Because decreasing α and increasing β increases the feasible set, the efficiency-fairness function is monotonically non-increasing with α for a fixed β and non-decreasing with β for a fixed α .

In the real world markets are not efficient. This is also the case in the telecommunications industry. There are many effects that distort market efficiency and lead to revenue functions that are non-smooth, non-continuous, non-concave and even non-monotonic. Clearly if a service provider has two customers one of which pays for the services and the other does not, the service provider will try to allocate more resources to serve the paying customer. However, if other considerations (social, regulations etc.) force the service provider to serve the non-paying customer, it will try to do it in a way that will maximize its revenues subject to certain “fairness” constraints. We provide here a framework that achieves this. In particular, we employ recently developed non linear programming (NLP) methods that can accurately solve the particular global optimization problems associated with maximizing efficiency subject to fairness constraints.

In this paper, we generalize our earlier work reported in [38] in two ways. First, we consider here a large class of non-linear (including non-concave, non-smooth, non-continuous and non-monotonic) utility functions, while [38] considered only linear utility functions. The use of non-concave, non-smooth and non-continuous utility

functions is motivated by the fact that markets are often inefficient. Another difference of [38] is that there we only considered the lower bound fairness parameter α , and here we also consider the upper bound parameter β . Such an upper bound is motivated as it avoids situations whereby a user is allocated significantly less bandwidth than his/her neighbor. In such a case, it will not be much consolation that the bandwidth the user allocated is not significantly lower than a certain overall “fair” value (the α constraint).

There have been many publications on fairness and other resource allocation problems in telecommunications networks (see for example [7,15,16,19,21,25,26,33,34] and references therein). Mathematically, the general approach is to maximize aggregate utility subject to linear capacity constraints. For tractability and to allow distributed flow control algorithms, the utility functions considered have usually been concave. Some formulations, such as [19], impose the additional constraint that each flow have a minimum and maximum transmission rate. The parameters α and β in our formulation provide a specific interpretation for these minimum and maximum rates.

The focus here is the problem of a network operator allocating virtual private links to users, rather than flow control. Instead of seeking to maximize the aggregate benefit to the users, this paper seeks to maximize the benefit to the service provider. More significantly, this problem allows centralized algorithms to be used, which allows a wider class of problems to be studied. By using global optimization algorithms, flows with non-concave, non-smooth and non-continuous utility functions can be considered.

Network designers must also choose the degree of unfairness allowed, by setting α and β . This can be done in terms of an efficiency-fairness function, which quantifies the tradeoff as follows. First, the “fair” rates are chosen, in terms of a fairness criterion such as max-min fairness or proportional fairness. Then, for a range of α and β values,

the operator's utility is numerically optimized given the (α, β) -fairness constraint, namely that each flow obtains between α and β times its fair allocation.

In this paper we will consider two global optimization algorithms. The first algorithm is the so-called Lipschitz Global Optimization (LGO). It is a well known algorithm that uses the Branch-and-bound global search method [28]. This algorithm is one of the best algorithms in solving optimization problems with constraints [27].

The second algorithm is the so-called Algorithm for Global Optimization Problems (AGOP) – a recently developed algorithm presented in [23]. This algorithm is designed for solving continuous optimization problems with box constraints, that is, problems where the feasible region is the Cartesian product of intervals. It is therefore relevant to the type of problems we consider in this paper to maximize efficiency subject to capacity and fairness constraints. The efficiency of the algorithm has been demonstrated in solving many difficult practical and test problems (see for example, [22], [23], [24]).

Throughout the paper we use the notation $\langle u, v \rangle$ for the link that connects nodes u and v and $[u, v]$ for the data flow from node u to node v .

As an illustration of the (α, β) -fairness concept, consider the two-node single-link network presented in Figure 1 with the link capacity equal to one. There are two flows from node u to node v , designated as Flow 1 and Flow 2. Both flows aim to transmit at unlimited rate. Assume that Flow 1 pays 1 [\$/unit capacity] and Flow 2 pays 1.5 [\$/unit capacity]. Assume that we choose max-min as our fairness criterion, and total revenue [\$] as our utility. Accordingly, if each of the flows is assigned a rate of $1/2$, this will yield a utility of 1.25. However, if we relax the fairness constraint to (α, β) -fairness, say for all $\{\alpha, \beta\}$ pairs that obey the relation $\beta = 1 - \ln \alpha$ for $1 \geq \alpha > 0$ (with $\beta = \infty$ for $\alpha = 0$), in which case the lower bound (set by α) is the

tighter constraint, then Flow 1 will be assigned the rate of $\alpha/2$ and Flow 2 will be assigned the rate of $1 - \alpha/2$. This will give utility of $1.5 - \alpha/4$ plotted in Figure 2. Maximum utility is achieved if fairness is completely ignored ($\alpha = 0$ and $\beta = \infty$) and Flow 2 is assigned the full bus capacity. Clearly, the slope of the efficiency-fairness




Fig. 1. Topology and flows of the two-node single-link example.




Fig. 2. The efficiency-fairness function for the two-node single-link example.

function, for the previous example, can be made significantly steeper if we further increase the charge of Flow 2 (and/or decrease the charge of Flow 1).

Figure 3 illustrates our framework. The inputs are network topology, a set of α , β values, efficiency utility function and fairness criterion, and the output is the efficiency fairness function for the set of the (α, β) values. Given the general setting of our framework, we can also answer questions of fairness associated with serving individuals or communities in remote locations. It is an important political and socio-economical problem in many countries how much society and telecommunications providers should spend in serving remote communities. One extreme view is that people in remote communities should have “equal access”. That is, they have the same access to telecommunications services, and at the same cost, people in major

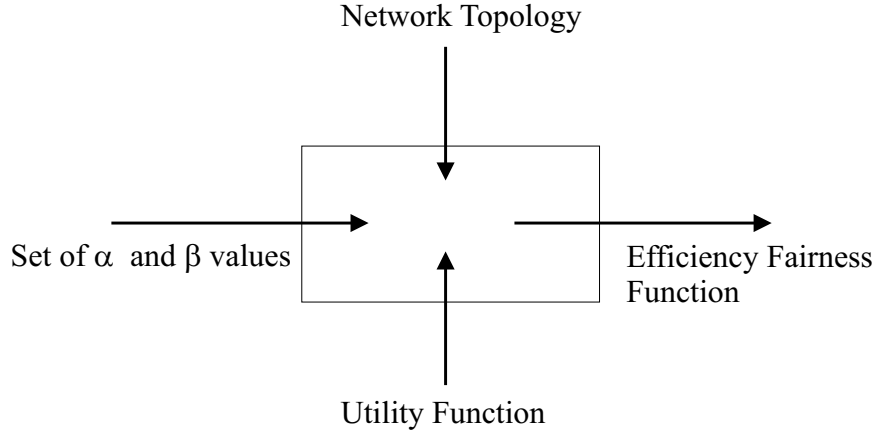


Fig. 3. The framework.

cities have. This corresponds in our framework to $\alpha = \beta = 1$. Another extreme view is that services to remote communities should be left to market forces $\alpha = 0$ and $\beta = \infty$. Of course there are many views in between these two extreme views. We will now use a three-node example to demonstrate how our framework can apply to the efficiency fairness tradeoff related to the question of servicing remote communities.

Consider the three-node network presented in Figure 4. Nodes A and B represent major cities. Node C is a remote community. Flow $[A, B]$ represents the flow from A to B on the directed link $\langle A, B \rangle$, and Flow $[B, A]$ represents the flow from B to A on a different directed link $\langle B, A \rangle$. The two links between nodes A and B have capacity 1. The link $\langle C, B \rangle$ does not yet exist in practice. There is a need to make a major investment in infrastructure to make this link a reality. To apply our framework to this problem, we include the link $\langle C, B \rangle$ with capacity bounded above by 1. The actual capacity of this link is to be determined, and this infrastructure investment will be included in the utility function associated with flow $[C, B]$; the cost of installing capacity is assumed to be twice the revenue raised over the life of the infrastructure. In other words, we assume for simplicity that allocating rate x_{CB} to flow $[C, B]$ will contribute to the aggregate utility $-2x_{CB} + x_{CB} = -x_{CB}$. This assumption of linear relationship between the rate provided and its infrastructure cost is made here

for simplicity. Normally, a more appropriate model for this relationship is a step function which we consider in Section 3. The fairness criterion we consider is equal rate fairness; that is, each flow is allocated rate 1. The aggregate utility of all three flows include the actual values of the flows $[A, B]$ and $[B, A]$ each of which obtains its maximum values at 1 and the utility of flow $[C, B]$ which is *minus* the flow on $[C, B]$. Altogether the utility is equal to $1 + 1 - x_{CB} = 2 - x_{CB}$. Notice that the fair allocation is 1 to each of the three flows, and that the capacity of each link is also equal to 1, the efficiency fairness curve does not change with β (no flow can exceed its fair allocation), so the allocation is independent of β . The rate obtained by flow $[C, B]$ will be α , and the efficiency-fairness function is $2 - \alpha$, shown in Figure 5. Of




Fig. 4. Topology and flows of the remote node example.




Fig. 5. The efficiency-fairness function for the remote node example.

course, the utility functions we considered above are simplistic. For example, serving customers at remote locations involves a much more complex cost function than the one we have considered. Nevertheless, the utility functions we consider in this paper have characteristics of very wide generality that can be applicable to realistic cases.

Having introduced the concepts of (α, β) -fairness, the efficiency-fairness function and having demonstrated their applicability, we are ready to formalize these notions. In Section 2 we formulate a nonlinear program (NLP) that leads to the efficiency-fairness function, and in Section 3 we further discuss this NLP and the methods used to solve it. In Section 4, we provide several network examples to demonstrate how to formulate an NLP that finds the optimal bandwidth allocation for a general network under the fairness and capacity constraints and how to obtain the efficiency-fairness function.

2 NLP Formulation

We consider an M node network. The nodes are designated $1, 2, 3, \dots, M$. All sources are assumed to be greedy. Each source may transmit one or more flows. Let R_{ij} be the rate assigned to flow $[i, j]$. The aim is to set the R_{ij} values to maximize the utility subject to fairness and capacity constraints.

Let F be the set of all flows. Consider the case that there is a utility function $U_{ij}(R_{ij})$ for flow $[i, j]$, and the total utility is the sum of the utility per flow,

$$U = \sum_{[i,j] \in F} U_{ij}(R_{ij}). \quad (1)$$

This form applies in many applications when users are charged according to their usage.

Let $f(i, j)$ be the fair allocation for flow $[i, j]$ according to our chosen fairness criterion. Let us assume that we require that the allocation is (α, β) -fair. Then, R_{ij} will be bounded below by $\alpha f(i, j)$ and above by $\beta f(i, j)$. This leads to the following nonlinear programming problem.

$$\text{Maximize } U = \sum_{[i,j] \in F} U_{ij}(R_{ij}) \quad (2)$$

$$\text{subject to: } \alpha f(i, j) \leq R_{ij} \leq \beta f(i, j), \quad \forall [i, j]; \quad (3)$$

$$\sum_{[i, j] \in F\langle m, n \rangle} R_{ij} \leq C\langle m, n \rangle, \quad \forall \langle m, n \rangle; \quad (4)$$

where $C\langle m, n \rangle$ denotes the capacity of link $\langle m, n \rangle$ (this may be the current capacity, or the maximum allowed value after a capacity upgrade) and $F\langle m, n \rangle$ is the set of all flows that use link $\langle m, n \rangle$.

3 Analysis and Solution Methods

For the sake of simplicity we will use the following notations. We denote the number of flows in the set $F = \{[i, j]\}$ by n . Let the variables x_k , $k = 1, 2, \dots, n$, stand for the rates R_{ij} assigned to flows $[i, j] \in F$. We also denote by m the number of all links in the network, and let C_l , $l \in \{1, 2, \dots, m\}$, be the capacity of link l . Let the usage of link l be specified as $r_l = [r_{l1}, r_{l2}, \dots, r_{ln}]$ where

$$r_{lk} = \begin{cases} 1 & \text{if flow } k \text{ uses link } l, \\ 0 & \text{otherwise.} \end{cases}$$

Then, given fairness parameters α and β , the problem (2)–(4) can be formulated as follows:

$$\text{maximize } U(x), \quad (5)$$

$$\text{s.t. } x \in X \subset \mathbb{R}_+^n; \quad (6)$$

where

$$X = \{x \geq 0 : x \in B(\alpha, \beta) \text{ and } r_l x \leq C_l, \quad l = 1, \dots, m\}, \quad (7)$$

and \mathbb{R}_+^n is the set of points in \mathbb{R}^n with non-negative coordinates. Throughout this paper, inequalities between vectors are taken component-wise.

The notation $r_l x$ stands for the scalar product of vectors r_l and x . The set $B(\alpha, \beta) \subset$

\mathbb{R}_+^n is a box defined by

$$B(\alpha, \beta) = \{x = (x_1, \dots, x_n) : \alpha f_k \leq x_k \leq \beta f_k, k = 1, \dots, n\},$$

where f_k is the fair allocation for flow k .

As mentioned above we consider $\alpha \in [0, 1]$ and $\beta \geq 1$. It is clear that, the feasible set X is a convex set (polytope). In this model, the fairness parameters α and β , are taken into account in box constraints.

In (5), $U(x)$ is a utility function of x representing what we call the efficiency.

In this paper, we will consider the class of efficiency (utility) functions defined by

$$U(x) = \sum_{k=1}^n \left(\max_{q=1, \dots, Q} a_k^q \ln(d_k x_k + b_k^q) - c_k S(x_k) \right), \quad (8)$$

where $a_k^q > 0$, $b_k^q > 0$, $c_k > 0$ and $d_k > 0$, for all $q = 1, \dots, Q$ and $k = 1, \dots, n$. The terms $c_k S(x_k)$ in (8) represent the costs of infrastructure upgrades. The function $S(x_k)$ may be discontinuous, such as a non-decreasing step function. Typically $S(0) = 0$. The framework does not require that these costs have the form of a sum over k , but that suffices for the network examples studied here.

Therefore, we consider the following Problem (**P1**):

$$\begin{aligned} \text{maximize} \quad & U(x) = \sum_{k=1}^n \left(\max_{q=1, \dots, Q} a_k^q \ln(d_k x_k + b_k^q) - c_k S(x_k) \right), \\ \text{s.t.} \quad & x \in X \subset \mathbb{R}_+^n; \end{aligned}$$

where X is defined by (7).

We say that x^* is a locally optimal solution (maximum) in Problem **P1** if there is a

neighborhood $x \in \mathcal{V}$ such that $U(y) \leq U(x^*)$ for all $y \in \mathcal{V} \cap X$.

Let us consider certain special cases with $c_k = 0$, $k = 1, \dots, n$. If $Q = 1$ and $a_k^1 = b_k^1 = 1$ for all k , then we have the following objective function:

$$U_1(x) = \sum_{k=1}^n \ln(d_k x_k + 1). \quad (9)$$

This is the simplest version of (8). In the calculation below we also consider the case when $Q > 1$. For example, if $Q = 2$ and $a_k^1 = 1$, $a_k^2 = 4$, $b_k^1 = 1$, $b_k^2 = 0.9$, and $d_k = 1$, then we have the following objective function

$$U_2(x) = \sum_{k=1}^n \max\{\ln(x_k + 1), 4 \ln(x_k + 0.9)\}. \quad (10)$$

Note that in the latter, $\ln(x_k + 1)$ dominates for $0 \leq x \leq 0.124$ and $4 \ln(x_k + 0.9)$ dominates for $0.124 < x$.

A few comments about the utility function (8) are in order.

Unlike the purely logarithmic utility function of [15,16,19,20], the utility function is well-behaved for all non-negative x because $b_k > 0$. In [15,16,19,20], the utility tends to $-\infty$ as a user's rate tends to zero, to ensure fairness. In our approach, a degree of fairness is imposed explicitly by requiring (α, β) -fairness, and so there is no need for an unbounded utility function. Bounded utility functions will generally better reflect the operator's true profit from allocating given rates.

The particular form of (8) is used here as an example of a utility function that is non-smooth and non-concave. In an inefficient market which is what we have in the real world, an operator may be able to charge users a higher price when the bandwidth becomes sufficient to provide a new service.

Now we note the following properties of Problem **P1**.

Proposition 1: If $Q = 1$ and $c_k = 0$ for all k , then Problem **P1** has a unique locally optimal solution.

The uniqueness of local maxima follows from the fact that the feasible set X is convex and the objective function is strictly concave because $a_k^1 > 0$ and $d_k^1 > 0$. Moreover, the condition $b_k^1 > 0$ guarantees that the objective function is well defined on the feasible set.

Proposition 2: If $c_k = 0$, then all locally optimal solutions to Problem **P1** belong to the boundary of the feasible set X .

This proposition follows from the fact that the feasible set X is convex and the objective function is monotonic because $c_k = 0$; that is, $U(x) \leq U(y)$ if $x \leq y$.

From Proposition 1 it follows that Problem **P1** with (9) has a unique locally optimal solution. However the situation is completely changes when $Q > 1$. In this case objective function $U(x)$ may be non-concave and non-smooth. Therefore it may have many local maxima, which are on the boundary of the feasible set. As an example, consider a simple case of two variables ($n = 2$) for objective function (10) with constraints $X = \{(x_1, x_2) \geq 0 : x_1 + x_2 \leq 1\}$. In this case we have 3 locally optimal solutions: $(0, 1)$, $(0.5, 0.5)$ and $(1, 0)$.

The number of locally optimal solutions may drastically increase as the number variables n increases. In this case the reasonable goal could be to find solutions that are close to the global solution.

The existence of many locally optimal solutions belonging to the boundary of the feasible set is the main difficulty that complicates the finding of a global solution to

Problem **P1**. This difficulty is similar to those in the concave minimization problem which is NP hard and is one of the challenging optimization problems (see for example [29]). Another difficulty is that the objective function (8) is non-smooth.

Finally, if $c_k \neq 0$, the objective function may be discontinuous. In most applications, S (and hence U) will be piecewise continuous. If S has a continuous pieces, then the box $B(\alpha, \beta)$ can be broken into a^n continuous sub-boxes. However, each sub-box may still contain multiple local minima.

These issues require the use of efficient (global) optimization techniques for solving Problem **P1**. However existing algorithms cannot in general guarantee to find a global solution. In many examples we know only “the best known solution”. The efficiency of an algorithm can be determined by comparison with other algorithms.

As mentioned above, we consider here the two global optimization algorithms LGO and AGOP. While LGO is well known, AGOP has only recently been developed. AGOP [23] is especially designed for continuous optimization problems with box constraints. It uses a line search mechanism where the descent direction is obtained via a dynamical systems approach. It is applicable to a wide range of optimization problems requiring only function evaluations to work. In particular it does not require gradient information and can be used to find minima of non-smooth functions. The efficiency of the algorithm has been demonstrated in solving many difficult practical and test problems (see for example, [22], [23], [24]).

The AGOP algorithm will now be described in terms of minimizing a cost function $g : \mathbb{R}^n \rightarrow \mathbb{R}$. AGOP must first be given a set of initial points, say $\Omega = \{x^1, \dots, x^q\} \subset \mathbb{R}^n$, $q \geq 2$. Generally, a suitable choice for an initial set of points is the set of some vertices of a given box. Let $x^* \in \Omega$ be the point in Ω with the smallest cost, that is, $g(x^*) \leq g(x)$ for all $x \in \Omega$. The set Ω and the values of g at each of the points in Ω

allow us to determine a vector v to be used as a possible descent direction from point x^* , as outlined below. An inexact line search along the direction of v provides a new point $\hat{x}^{q+1} \neq x^*$. A local search about \hat{x}^{q+1} is then carried out using a method called *local variation*. This is an efficient local optimization technique that does not explicitly use derivatives and can be applied to non-smooth functions. A good survey of direct search methods can be found in [17]. Letting x^{q+1} denote the optimal solution of this local search, the set Ω is augmented to include x^{q+1} . Starting with this updated Ω , the whole process can be repeated. The process is terminated when v is approximately 0 (or a prescribed bound on the number of iterations is reached). The solution returned is the current x^* , that is, the point in Ω with the smallest cost.

The success of global optimization algorithms mainly depends on their ability to escape the best local minimum found so far in order to find “deeper” local minima. In our case this is the determination of a possible descent direction v from currently found local minimum x^* . We present here the formula used in the calculations below. For more details and motivations behind it see [20].

Let $I = \{1, \dots, q\}$ and let $\Omega = \{x^m : m \in I\}$ be a set of initial points, and $g^m = g(x^m)$. Let $g^* = g(x^*)$, with $x^* = (x_1^*, \dots, x_n^*)$, be the smallest cost of the points in Ω . For each coordinate $i \in \{1, \dots, n\}$, define the following sets: $X_i^+ = \{m \in I : x_i^m > x_i^*\}$, $X_i^- = \{m \in I : x_i^m < x_i^*\}$, $G^+ = \{m \in I : g^m > g^*\}$, $G_i^{++} = G^+ \cap X_i^+$, $G_i^{+-} = G^+ \cap X_i^-$. Let $|A|$ denote the cardinality of a set A .

Define $v = (v_1, \dots, v_n)$ as follows:

$$\begin{aligned}
v_i = & \sum_{m \in G_i^{+-}} \frac{1}{|G^+|} \frac{-\Delta x_i^m}{\Delta g^m} \alpha_i^m \cdot \sum_{m \in G_i^{+-}} \frac{1}{|X_i^-|} \frac{\Delta g^m}{\|\Delta x^m\|} \alpha_i^m \\
& - \sum_{m \in G_i^{++}} \frac{1}{|G^+|} \frac{\Delta x_i^m}{\Delta g^m} \alpha_i^m \cdot \sum_{m \in G_i^{++}} \frac{1}{|X_i^+|} \frac{\Delta g^m}{\|\Delta x^m\|} \alpha_i^m.
\end{aligned} \tag{11}$$

Here $\Delta x_i^m = x_i^m - x_i^*$, $\Delta g^m = g^m - g^*$, $\alpha_i^m = |\Delta x_i^m|/\|\Delta x^m\|$ and $\|\Delta x^m\|$ is the Euclidean norm.

Formula (11) represents a deterministic approach to calculate v , in contrast to randomized global search algorithms. It tries to take into account the contribution of each coordinate i on the increase in function values. The final value v_i is, in some sense, the average value over the all set Ω .

We consider v as a possible “global” descent direction because it uses information obtained from points that may be quite far from each other. If this direction fails to provide a better point x^{q+1} , that is $g(x^{q+1}) > g(x^*)$, then we add this point to set Ω , which supplies a new direction v calculated by formula (11) with this updated set. We note that, the role of a direction v is to find a basin of a new local minimum; that is to find \hat{x}^{q+1} rather than x^{q+1} . In other words, a search over direction v is successful if it finds a basin of a deeper local minima.

In global optimization, usually, we can not guarantee that we will find a successful descent direction even if x^* is not a global minima. This is the main difficulty that global optimization algorithms encounter. The success of an algorithm can be checked only on numerical experiments. The results obtained in this paper once more emphasize the efficacy of the proposed approach.

If we have just one initial point x^1 , to run this algorithm starting from this particular point, first we apply a local search about this point to get a new point (local minimum) x^2 . Then, the above procedure is performed using the set $\Omega = \{x^1, x^2\}$.

To solve optimization problems with constraints, AGOP uses the following scheme, that is demonstrated on Problem **P1**.

First, we transform this problem to a minimization problem with box constraints,

applying penalty functions to the linear constraints (see [30] and references therein).

Given a penalty coefficient $\gamma > 0$, set

$$g_\gamma(x) = -U(x) + \gamma \sum_{l=1}^m (\max\{r_l x - C_l, 0\})^2. \quad (12)$$

Consider the Problem (**P2**):

$$\text{Minimize } g_\gamma(x), \quad \text{s.t. } x \in B(\alpha, \beta).$$

Then we perform the following steps to solve Problem **P1**.

Step 1. Take any penalty coefficient $\gamma > 0$ and any number $\lambda > 1$.

Step 2. Apply AGOP to Problem **P2**, which has only box constraints and denote the solution by x^0 .

Step 3. Set $p = 1$ and $x^p = x^0$.

Step 4. Set $\gamma = \gamma\lambda$. Apply AGOP to **P2** starting from initial point x^p . Let the solution found be x^{p+1} .

Step 5. If $|g_\gamma(x^{p+1}) - g_\gamma(x^p)| < \varepsilon$ then stop. Otherwise set $p = p + 1$ and go to Step 4.

The convergence of the algorithm to a local minimum can be proved for smooth functions, in particular, for Problem **P1** with (9). AGOP, like any other algorithms using direct search methods [17], does not guarantee to get a local minimum for non-smooth objective functions. However, in practice, these methods (having global search character) often perform well when dealing with non-smooth functions, (see for example [22]), and also even non-continuous functions. The results obtained for Examples 2 and 3 in Section 4 also confirm this fact (see Tables 2, 4 for a comparison of LGO and AGOP on non-smooth objective function U_2).

4 Numerical Results

The flexibility of the form of the objective function U in **P1** means that powerful global optimization techniques are required to perform the design. This section uses two such techniques, LGO and AGOP, to determine the efficiency-fairness curve for a range of topologies, with objective functions of the form (8).

The (α, β) -fairness concept specifies how far a rate allocation can deviate from the “fair” rate, but does not specify what rate is to be used as the reference. For these results, we take the max-min fair rates [2] to be “fair”. Given α and β , this yields box constraints on the rates, of the form

$$B(\alpha, \beta) = \{x = (x_1, \dots, x_n) : \alpha f_k \leq x_k \leq \beta f_k, k = 1, \dots, n\},$$

where f_k is the fair allocation for flow k . Given box $B(\alpha, \beta)$, we will denote by $\xi(\alpha, \beta)$ the optimal (maximal) value of objective function in Problem **P1**:

$$\xi(\alpha, \beta) = \max_{x \in X} U(x).$$

In each of the examples below, the set X represents the set of feasible points for the example under consideration.

The efficiency-fairness tradeoff is governed by the values of α and β . In order to make the efficiency-fairness function easier to plot, we express both α and β in terms of a single parameter, s . In particular, we set

$$\alpha_s = (10 - s)/10, \quad \beta_s = (\sqrt{2})^s, \quad s = 0, 1, \dots, 10.$$

It is clear that

$$B(\alpha_0, \beta_0) \subset B(\alpha_1, \beta_1) \subset \cdots \subset B(\alpha_{10}, \beta_{10});$$

and therefore

$$\xi(\alpha_0, \beta_0) \leq \xi(\alpha_1, \beta_1) \leq \cdots \leq \xi(\alpha_{10}, \beta_{10}).$$

In all examples below, we will consider objective functions given by (8) with $Q = 1$ or $Q = 2$.

4.1 Example 1: Linear Network with uniform capacity.

Consider the network shown in Fig. 1 of [34]. There are n flows and $(n - 1)$ concatenated links in the network. We consider the case that each link has bandwidth of 1 unit. Flow x_n travels through all the links; flow x_1 travels through link 1; flow x_2 travels through link 2; and so on, and flow x_{n-1} travels through link $n - 1$, as shown in Figure 6.

The Max-Min fair allocation for this network is to assign the same rate to all the flows:

$$x_1 = x_2 = x_3 = \cdots x_n = \frac{1}{2}$$

where x_i is the rate allocation to flow i . We describe the optimizing problem of efficiency-fairness tradeoff as follows:

$$\begin{aligned} & \text{Maximize} && U(x) \\ & \text{subject to:} && \\ & \frac{1}{2}\beta \geq x_i \geq \frac{1}{2}\alpha && \text{for } i = 1, 2, \dots, n \\ & x_i + x_n \leq 1 && \text{for } i = 1, 2, \dots, n - 1 \\ & x_i \geq 0 && \text{for } i = 1, 2, \dots, n \end{aligned}$$

For this example, we consider the following two objective functions:




Fig. 6. The classical network model used in Example 1.

$$U_{11}(x) = \sum_{i=1}^n \ln(x_i + 1),$$

and

$$U_{12}(x) = \sum_{i=1}^n \max\{\ln(x_i + 1), 4 \ln(x_i + 0.9)\}.$$

Given α and β , the maximal values of objective functions will be denoted as:

$$\xi_{11}(\alpha, \beta) = \max U_{11}(x),$$

and

$$\xi_{12}(\alpha, \beta) = \max U_{12}(x),$$

respectively.

It is not difficult to show that objective functions U_{11} and U_{12} have a unique optimal solution for this simple example. Both algorithms, LGO and AGOP could easily find the optimal solution for all boxes $B_s = B(\alpha_s, \beta_s)$, $s = 0, \dots, 10$. The results are presented in Table 1. The efficiency-fairness functions for Example 1 for the objective functions $U_{11}(x)$ and $U_{12}(x)$ are presented in Figure 7.

4.2 Example 2: Linear Network with two long flows.

Consider now a somewhat more complicated network — the linear network with two long flows [34], as shown in Figure 8. In this model, we choose the capacity

Table 1

The results obtained by algorithms LGO and AGOP for objective functions U_{11} , U_{12} with $n = 20$. Here $B_s = B(\alpha_s, \beta_s)$ ($s = 0, 1, \dots, 10$) and ξ_{11} , ξ_{12} are the optimal values of objective functions. Note that the results are the same for both algorithms.

Box	ξ_{11}	ξ_{12}
B_0	8.1093	26.9178
B_1	8.6984	29.4393
B_2	9.2665	31.8648
B_3	9.8148	34.1999
B_4	10.3443	36.4496
B_5	10.8558	38.6180
B_6	11.3503	40.7090
B_7	11.8283	42.7260
B_8	12.2905	44.7671
B_9	12.7376	46.8029
B_{10}	13.1698	48.7809




Fig. 7. Efficiency-fairness functions for Example 1.

$C = (500, 400, 300, 200, 500)^T$, and calculate the max-min fairness rate allocation. It is given by: $x_1 = 400$, $x_2 = 300$, $x_3 = 100$, $x_4 = 100$, $x_5 = 400$, $x_6 = 100$, $x_7 = 100$.

The optimization problem is as follows:




Fig. 8. The Linear network with two long flows used in Example 2.

$$\text{Maximize } U(x)$$

subject to:

$$\begin{aligned} 400\beta &\geq x_1 \geq 400\alpha, & 300\beta &\geq x_2 \geq 300\alpha, \\ 100\beta &\geq x_3 \geq 100\alpha, & 100\beta &\geq x_4 \geq 100\alpha, \\ 400\beta &\geq x_5 \geq 400\alpha, & 100\beta &\geq x_6 \geq 100\alpha, \\ 100\beta &\geq x_7 \geq 100\alpha, & x_1 + x_7 &\leq 500, \\ x_2 + x_7 &\leq 400, & x_3 + x_6 + x_7 &\leq 300, \\ x_4 + x_6 &\leq 200, & x_5 + x_6 &\leq 500, \\ x_i &\geq 0 \text{ for } i = 1, 2, \dots, 7. \end{aligned}$$

For this example, we consider the following two utility functions:

$$U_{21}(x) = \sum_{i=1}^7 \ln(x_i/500 + 1)$$

and

$$U_{22}(x) = \sum_{i=1}^5 \ln(x_i/500 + 1) + \sum_{i=6}^7 \max\{\ln(x_i/500 + 1), 4 \ln(x_i/500 + 0.9)\}.$$

These two utility functions will lead to the efficiency-fairness functions:

$$\xi_{21}(\alpha, \beta) = \max_{x \in X} U_{21}(x)$$

and

$$\xi_{22}(\alpha, \beta) = \max_{x \in X} U_{22}(x),$$

respectively.

The results are presented in Table 2 and Figure 9. As mentioned above, objective function U_{21} has a unique optimal solution. Both algorithms, LGO and AGOP could easily find this solution for all boxes $B(\alpha_s, \beta_s)$, $s = 0, \dots, 10$. However, objective function U_{22} has many locally optimal solutions. As can be seen from Figure 9, the LGO algorithm has a difficulty in finding deep locally optimal solutions when the box $B(\alpha_s, \beta_s)$ becomes larger (for larger s values). The estimate from LGO can be improved slightly by enforcing the monotonicity constraint, increasing the estimates for certain (α, β) pairs as necessary; this is not necessary using AGOP.

To study the effect of the α and β parameters of the efficiency-fairness function we used only AGOP to produce the results presented in Figures 10 and 11.

The results show that an appreciable gain in utility can be achieved by relaxing the fairness constraints.

Note that in Figures 10 and 11, the curves are flat for a wide range of α and β values. The operating point, and hence utility, becomes independent of β when the β constraint ceases to be tight; that is, when all rates are constrained from decreasing by the α constraints, or constrained from increasing by the capacity constraints. Clearly, the further α is from 1, the larger the utility is at the point where this happens.

Table 2

The results obtained by algorithms LGO and AGOP for objective functions U_{21} and U_{22} . Here $B_s = B(\alpha_s, \beta_s)$ ($s = 0, 1, \dots, 10$) and ξ_{21} , ξ_{22} are the optimal values of objective functions. Note that the results for U_{21} are the same for both algorithms.

Box	ξ_{21}	ξ_{22}	
	LGO and AGOP	LGO	AGOP
B_0	2.3749	2.7727	2.7727
B_1	2.4251	2.8052	2.8052
B_2	2.4730	2.8340	2.8340
B_3	2.5188	2.8618	2.8617
B_4	2.5623	2.8886	2.8883
B_5	2.6038	2.8053	2.9508
B_6	2.6432	2.8451	3.0067
B_7	2.6805	2.6805	3.0566
B_8	2.7158	2.7158	3.1004
B_9	2.7492	2.7492	3.1383
B_{10}	2.7806	2.7806	3.1703

An analogous effect is observed in Figure 11. Note however that the curves for different β values are not entirely coincident until they reach the β -dependent ceiling as α decreases. Rather there is a gradual divergence before the limit is reached. That occurs when some of the flows' rates are constrained by α and some by β .

4.3 Example 3: A 12-node Network

The topology of this network is presented in the Figure 12. Assume that the capacity of each link is set at $C = 2$, and that there are 35 flows in this model. The flows' route information and the max-min rate allocations are given in Table 3. Considering max-min fairness and a general utility function, the relevant NLP is thus formulated as follows:




Fig. 9. Efficiency-fairness functions for Example 2 with objective functions $U_{21}(x)$ and $U_{22}(x)$.




Fig. 10. Efficiency-fairness functions for Example 2, where α is fixed and β changes.

Maximize $U(x)$;

subject to:

$$(0 \leq) \alpha x_i^\dagger \leq x_i \leq \beta x_i^\dagger \quad i = 1, \dots, 35$$

$$x_1 + x_2 + x_4 + x_5 + x_9 + x_{13} + x_{33} + x_{35} \leq C,$$

$$x_3 + x_6 + x_7 + x_9 + x_{13} + x_{34} \leq C,$$

$$x_2 + x_8 + x_{12} + x_{13} + x_{14} + x_{17} + x_{35} \leq C,$$

$$x_4 + x_5 + x_{10} + x_{11} + x_{14} + x_{17} + x_{33} \leq C,$$

$$x_{12} + x_{15} + x_{16} + x_{35} \leq C, \quad x_6 + x_7 + x_{34} \leq C,$$

$$x_{11} + x_{18} + x_{19} + x_{24} + x_{26} + x_{27} + x_{30} \leq C,$$

$$x_5 + x_{17} + x_{20} + x_{21} + x_{22} + x_{23} + x_{26} + x_{27} + x_{30} + x_{33} \leq C,$$

$$x_7 + x_{15} + x_{19} + x_{24} + x_{25} + x_{28} + x_{29} + x_{30} + x_{34} \leq C,$$

$$x_{24} + x_{28} + x_{29} + x_{31} + x_{32} + x_{34} + x_{35} \leq C, \quad x_{23} \leq C,$$

$$x_{21} + x_{22} + x_{23} + x_{27} + x_{33} \leq C, \quad x_{22} + x_{33} \leq C,$$

$$x_{28} \leq C, \quad x_{31} + x_{35} \leq C.$$




Fig. 11. Efficiency-fairness functions for Example 2, where β is fixed and α changes.




Fig. 12. The 12 node network used in Example 3.

For this example, we consider the following two utility functions:

$$U_{31}(x) = \sum_{i=1}^{35} \ln(x_i + 1);$$

and

$$U_{32}(x) = \sum_{i \notin I} \ln(x_i + 1) + \sum_{i \in I} \max\{\ln(x_i + 1), 4 \ln(x_i + 0.9)\},$$

where the set

$$I = \{5, 13, 17, 24, 26, 27, 30, 33, 34, 35\}$$

represents routes on which there is demand for a “premium” service, if the rate is sufficient.

Table 3
Routes and calculated max-min rate allocation of the 12 node network.

Flows (O-D)	Route	Max-min Rate Allocation (c)
1 - 2	1 - 2	$x_1^\dagger = 2/15$
1 - 3	1 - 2 - 3	$x_2^\dagger = 2/15$
1 - 4	1 - 4	$x_3^\dagger = 61/240$
1 - 5	1 - 2 - 5	$x_4^\dagger = 2/15$
1 - 8	1 - 2 - 5 - 8	$x_5^\dagger = 1/10$
1 - 6	1 - 4 - 6	$x_6^\dagger = 61/240$
1 - 7	1 - 4 - 6 - 7	$x_7^\dagger = 9/80$
2 - 3	2 - 3	$x_8^\dagger = 1/6$
2 - 4	2 - 1 - 4	$x_9^\dagger = 2/15$
2 - 5	2 - 5	$x_{10}^\dagger = 1/5$
2 - 6	2 - 5 - 6	$x_{11}^\dagger = 1/5$
2 - 7	2 - 3 - 7	$x_{12}^\dagger = 1/6$
3 - 4	3 - 2 - 1 - 4	$x_{13}^\dagger = 2/15$
3 - 5	3 - 2 - 5	$x_{14}^\dagger = 1/6$
3 - 6	3 - 7 - 6	$x_{15}^\dagger = 9/80$
3 - 7	3 - 7	$x_{16}^\dagger = 2/15$
3 - 8	3 - 2 - 5 - 8	$x_{17}^\dagger = 1/10$
5 - 6	5 - 6	$x_{18}^\dagger = 11/40$
5 - 7	5 - 6 - 7	$x_{19}^\dagger = 9/80$
5 - 8	5 - 8	$x_{20}^\dagger = 1/10$
5 - 9	5 - 8 - 9	$x_{21}^\dagger = 1/10$
5 - 10	5 - 8 - 9 - 10	$x_{22}^\dagger = 1/10$
5 - 11	5 - 8 - 9 - 11	$x_{23}^\dagger = 1/10$
5 - 12	5 - 6 - 7 - 12	$x_{24}^\dagger = 9/80$
6 - 7	6 - 7	$x_{25}^\dagger = 9/80$
6 - 8	6 - 5 - 8	$x_{26}^\dagger = 1/10$
6 - 9	6 - 5 - 8 - 9	$x_{27}^\dagger = 1/10$
6 - 10	6 - 7 - 12 - 10	$x_{28}^\dagger = 9/80$
6 - 12	6 - 7 - 12	$x_{29}^\dagger = 9/80$
7 - 8	7 - 6 - 5 - 8	$x_{30}^\dagger = 1/10$
7 - 11	7 - 12 - 11	$x_{31}^\dagger = 5/24$
7 - 12	5 - 12	$x_{32}^\dagger = 5/24$
1 - 10	1 - 2 - 5 - 8 - 9 - 10	$x_{33}^\dagger = 1/10$
1 - 12	1 - 4 - 6 - 7 - 12	$x_{34}^\dagger = 9/80$
1 - 11	1 - 2 - 3 - 7 - 12 - 11	$x_{35}^\dagger = 1/10$

These two utility functions will lead to the efficiency-fairness functions:

$$\xi_{31}(\alpha, \beta) = \max_{x \in X} U_{31}(x)$$

and

$$\xi_{32}(\alpha, \beta) = \max_{x \in X} U_{32}(x),$$

respectively.

The results are presented in Table 4. The performance of algorithms LGO and AGOP

Table 4

The results obtained by algorithms LGO and AGOP for objective functions U_{31} and U_{32} . Here $B_s = B(\alpha_s, \beta_s)$ ($s = 0, 1, \dots, 10$) and ξ_{31} , ξ_{32} are the optimal values of objective functions. Note that the results for U_{31} are the same for both algorithms.

Box	ξ_{31}	ξ_{32}	
		LGO	AGOP
B_0	4.5537	4.5597	4.5597
B_1	6.2648	6.8061	6.8061
B_2	8.5403	10.8493	10.8493
B_3	9.5563	13.0041	13.0041
B_4	10.2515	14.2254	14.1813
B_5	10.6974	14.9410	14.8482
B_6	10.9704	15.0653	15.0728
B_7	11.1267	15.0614	15.2953
B_8	11.1747	14.4944	15.5612
B_9	11.2196	13.8754	15.7534
B_{10}	11.2624	13.8102	15.8245

are similar to those in Example 2. The objective function U_{31} has a unique optimal solution and both algorithms finds this solution for all boxes $B_s = B(\alpha_s, \beta_s)$, $s = 0, \dots, 10$. However, objective function U_{32} has many locally optimal solutions. LGO provides slightly better solutions for boxes B_4 and B_5 , however, AGOP performs better on larger boxes, and provides a monotonic estimate of the efficiency-fairness curve.

LGO was run in the General Algebraic Modeling System (GAMS) and AGOP in Lahey Fortran. The aggregate elapsed times for all 11 cases ($s = 0, \dots, 10$) was around 30% more for AGOP than LGO.

4.4 Example 4: A case with a remote node

We now consider an example with the four nodes network shown in Figure 14, where $c_l = 1$ for all links. In this example, nodes 2, 3 and 4 are three cities with existing




Fig. 13. Efficiency-fairness functions for the 12-node network of Example 3.

telecommunications infrastructure and node 1 represents a remote location. The link from node 1 to node 2 does not exist and a major cost is required to connect node 1 to the rest. We consider the following four flows: $x_1 = [1, 4]$, $x_2 = [2, 3]$, $x_3 = [3, 4]$ and $x_4 = [4, 2]$. Each of these flows is an aggregate of many individual flows. Let us assume that the three flows x_2, x_3, x_4 are equal in size, that is, each of them has the same number of individual flows. However, flow x_1 that is associated with traffic from the remote node 1, is much smaller than the other three. In this example, we consider the case that the number of individual flows that it carries is 19 times smaller than that of each of the other three flows. In this case, we will use the weighted Max-Min fairness criterion to reflect the difference in the number of individual flows carried by x_2, x_3, x_4 versus x_1 in a way that individual users will obtain “equal access”. Accordingly, the weighted Max-Min rates are:

$$x_1 = \frac{1}{20}, \quad x_2 = x_3 = \frac{19}{20}, \quad x_4 = 1.$$

Considering this fairness criterion, relevant NLP is formulated as follows:




Fig. 14. Topology and routes for the case with a remote node used in Example 4.

$$\begin{aligned}
 & \text{Maximize} && U(x) \\
 & \text{subject to:} \\
 & \frac{1}{20}\beta \geq x_1 \geq \frac{1}{20}\alpha, \\
 & \beta \geq x_4 \geq \alpha \\
 & \frac{19}{20}\beta \geq x_i \geq \frac{19}{20}\alpha \quad \text{for } i = 2, 3, \\
 & x_1 + x_2 \leq 1 \\
 & x_1 + x_3 \leq 1 \\
 & x_4 \leq 1 \\
 & x_i \geq 0 \quad \text{for } i = 1, 2, 3, 4.
 \end{aligned}$$

For this example, we consider the objection functions U_{40} , U_{41} and U_{42} , where

$$U_{4k}(x) = -S(x_1) + 2^k \ln(x_1 + 1) + \sum_{k=2}^4 \ln(x_k + 1),$$

$S(x) = 0$ if $x = 0$, $S(x) = 1$ if $x \in (0, 1/3]$, $S(x) = 1.5$ if $x \in (1/3, 2/3]$, and $S(x) = 2$ if $x > 2/3$. These utility functions lead to the efficiency-fairness functions:

$$\xi_{40}(\alpha, \beta) = \max_{x \in X(\alpha, \beta)} U_{40}(x), \quad \xi_{41}(\alpha, \beta) = \max_{x \in X(\alpha, \beta)} U_{41}(x),$$

and

$$\xi_{42}(\alpha, \beta) = \max_{x \in X(\alpha, \beta)} U_{42}(x),$$

respectively.

We used AGOP to solve this problem. The results are presented in Table 5. The

Table 5

The results obtained by algorithm AGOP for objective functions U_{40} , U_{41} , U_{42} . Here $B_s = B(\alpha_s, \beta_s)$ ($s = 0, 1, \dots, 10$) and ξ_{40} , ξ_{41} , ξ_{42} are the optimal values of objective functions.

Box	ξ_{40}	ξ_{41}	ξ_{42}
B_0	1.0776	1.1264	1.2240
B_1	1.0779	1.1441	1.2807
B_2	1.0783	1.1675	1.3581
B_3	1.0785	1.1973	1.4619
B_4	1.0788	1.2334	1.5980
B_5	1.0790	1.2726	1.7708
B_6	1.0791	1.2902	1.8655
B_7	1.0793	1.2902	1.8655
B_8	1.0794	1.2902	1.8655
B_9	1.0794	1.2902	1.8655
B_{10}	2.0794	2.0794	2.0794

efficiency-fairness functions $\xi_{4i}(\alpha, \beta)$, $i = 0, 1, 2$, are presented in Figure 15. Note




Fig. 15. Efficiency-fairness functions for Example 4.

that when $\alpha = 0$, it is permissible to set $x_1 = 0$ and not to build the new link at all. For the three objective functions considered, this is the optimal strategy in that case,

giving $x^* = (0, 1, 1, 1)$. When $x_1 = 0$ is not allowed, there is little to be gained by relaxing the fairness requirements if the utility function is U_{40} or U_{41} , but significantly higher revenue can be obtained by sacrificing fairness if the utility has the form of U_{42} .

5 Conclusions

This paper has presented the (α, β) -fairness concept, which specifies how far a resource allocation is allowed to be from an “ideally fair” allocation, according to any fairness criterion. This decouples the efficiency criterion from the fairness criterion and quantifies the familiar concept of the tradeoff between efficiency-fairness.

We have described a framework to produce “efficiency-fairness functions” that allow network operators to first set fairness constraints and then to optimize their efficiency. We have formulated an NLP problem which finds the optimal rate allocation for a general network and any ideally fair rate allocation, under the (α, β) -fairness constraints. This leads to what we call the efficiency-fairness function, which shows the increase in efficiency as a function of the extent to which fairness is compromised.

This framework applies both when the network is fixed and only the rates can be chosen, and also when new capacity is to be added to the network. In the latter case, the “efficiency” reflects both the revenue and the infrastructure cost associated with the rate allocation.

We have applied two global optimization algorithms, LGO and AGOP, to solving the NLP for a variety of networks. For the examples tested, AGOP seems to be a particularly promising algorithm.

Acknowledgement

The authors thank Alex Rubinov for reviewing the manuscript and helpful comments.

References

- [1] Bazaraa MS, Jarvis JJ. Linear programming and network flows. John Wiley & Sons, ISBN: 0-471-06015-1; 1976.
- [2] Bertsekas D, Gallager R. Data networks. Prentice Hall; 1992.
- [3] Bolton GE, Ockenfels A. How do efficiency and equity trade-off when a majority rules? Working paper, University of Cologne 2003, online; available: <http://ockenfels.uni-koeln.de/download/papers/fair-vs-eff.pdf>
- [4] Bonald T, Massoulié L. Impact of fairness on Internet performance. In: Proceedings of ACM Sigmetrics'01, June 2001.p. 82-91.
- [5] Butler M, Williams HP. Fairness versus efficiency in charging for the use of common facilities. Journal of Operational Research Society 2002;53(12): 1324-329.
- [6] Charny A, Clark DD, Jain R. Congestion control with explicit rate indication. In: Proceedings of IEEE International Conference on Communications(ICC95), Seattle; June 1995.p. 18-22.
- [7] Cui T, Andrew LLH, Zukerman M, Tan L. Improving the Fairness of FAST TCP to New Flows. IEEE Communications Letters 2006;10(5): 414-416.
- [8] Davik F, Yilmaz M, Gjessing S, Uzun N. IEEE 802.17 resilient packet ring tutorial. IEEE Communications Magazine 2004;42(3): 112-118.
- [9] Gambiroza V, Liu Y, Yuan P, Knightly E. High performance fair bandwidth allocation for resilient packet rings. In: Proceedings of the 15th ITC Specialist Seminar on Traffic Engineering and Traffic Management, Wurzburg, Germany; July 2002.p. 22-24.
- [10] Gambiroza V, Yuan P, Balzano L, Liu Y, Sheafor S, Knightly E. Design, analysis, and implementation of DVSR: A fair high-performance protocol for packet rings. IEEE/ACM Transactions on Networking 2004;12(1): 85-102.
- [11] Güth W, Kliemt H, Ockenfels A. Fairness versus efficiency: An experimental study of (mutual) gift giving. Journal of Economic Behaviour and Organization 2003;50: 465-475.
- [12] Hahne EL, Choudhury AK, Maxemchuk NF. DQDB networks with and without bandwidth balancing. IEEE Transactions on Communications 1992;40(7): 1192-1204.
- [13] Hawkins CN, Green J, Sharma M, Vasani K. Resilient packet rings for metro networks. August 2001. <http://www.rpralliance.org/>.
- [14] Hsu SL. Fairness versus efficiency in environmental law. Ecology Law Quarterly 2004;31: 303-401.
- [15] Kelly F. Charging and rate control for elastic traffic. European Transactions on Telecommunications 1997;8: 33-37.
- [16] Kelly F, Maulloo A, Tan D. Rate control for communication networks: Shadow prices, proportional fairness and stability. Journal of Operations Research Society 1998;49(3): 237-252.

- [17] Kolda TG, Lewis RM, Torczon V. Optimization by direct search: New perspectives on some classical and modern methods. *SIAM Review* 2003;45(3): 385-482.
- [18] Kunniyur S, Srikant R. End-to-end congestion control: utility functions, random losses and ECN marks. *IEEE/ACM Transactions on Networking* 2003;11(5): 689-702.
- [19] Low SH, Lapsley DE. Optimization flow control I: Basic algorithm and convergence. *IEEE/ACM Transactions on Networking* 1999;7: 861-875.
- [20] Low SH, Peterson L, Wang L. Understanding TCP Vegas: A duality model. *Journal of the ACM* 2002;49(2): 207-235.
- [21] Luo H, Lu S, Bharghavan V, Cheng J, Zhong G. A packet scheduling approach to QoS support in multihop wireless networks. *ACM Journal of Mobile Networks and Applications (MONET)*, Special Issue on QoS in Heterogeneous Wireless Networks 2004;9(3): 193-206.
- [22] Mammadov MA, Orsi R. H_∞ synthesis via a nonsmooth, nonconvex optimization approach. *Pacific Journal of Optimization* 2005;1(2): 405-420.
- [23] Mammadov MA, Rubinov AM, Yearwood J. Dynamical systems described by relational elasticities with applications to global optimization. In V. Jeyakumar and A. Rubinov (Eds.) *Continuous optimisation: current trends and applications*. Springer; 2005.p. 365-387.
- [24] Mammadov MA, Yearwood J, Aliyeva L. Multi label classification and drug-reaction associations using global optimization techniques. In: *Proceedings of The Sixth International Conference on Optimization: Techniques and Applications (ICOTA6)*, Ballarat; December 2004.
- [25] Massoulié L, Roberts J. Bandwidth sharing: Objectives and algorithms. In: *Proceedings of INFOCOM'99*, New York; March 1999.p. 1395-1403.
- [26] Mo J, Walrand J. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking* 2000;8(5): 556-567.
- [27] Neumaier A, Shcherbina O, Huyer W, Vinko T. A comparison of complete global optimization solvers, [online]. Available: <http://citeseer.ifi.unizh.ch/642030.html>.
- [28] Pinter J. *LGO - A model development system for continuous global optimization; User's Guide (Revised edition)*, Pinter Consulting Services, Halifax, NS Canada, 2003.
- [29] Pardalos PM, Rosen JB. Methods for global concave minimization: A bibliographic survey. *SIAM Review* 1986;28(3): 367-379.
- [30] Rubinov AM, Yang X. *Lagrange-Type Functions in Constrained Non-Convex Optimization*. Kluwer Academic; 2003.
- [31] Srinivasan R, Somani A. On achieving fairness and efficiency in high-speed shared medium access. *IEEE/ACM Transactions on Networking* 2003;11(1): 111-124.
- [32] Tan L, Pugh AC, Yin M. Rate-based congestion control in ATM switching networks using a recursive digital filter. *Control Engineering Practice* 2003;11(10): 1171-1181.

- [33] Tan L, Zhang X, Andrew LLH, Chan S, Zukerman M. Price-based max-min fair rate allocation in wireless multi-hop networks. *IEEE Communications Letters* 2006;10(1): 31-33.
- [34] Tang A, Wang J, Low SH. Is fair allocation always inefficient? In: *Proceedings of IEEE INFOCOM*, Hong Kong; 2004.p. 35-45.
- [35] Vandalore B, Fahmy S, Jain R, Goyal R, Goyal M. A definition of general weighted fairness and its support in explicit rate switch algorithms. In: *Proceeding of Sixth International Conference on Network Protocols (ICNP)*, Austin, Texas, USA; October 1998.p. 22-30.
- [36] Vandalore B, Fahmy S, Jain R, Goyal R, Goyal M. General weighted fairness and its support in explicit rate switch algorithms. *Computer Communications* 2000;23(2): 149-161.
- [37] Zukerman M, Potter P. The DQDB protocol and its performance under overload traffic conditions. *Computer Network and ISDN Systems* 1990;20(1-5): 261-270.
- [38] Zukerman M, Tan L, Wang H, Ouveysi I. Efficiency-fairness tradeoff in telecommunications networks. *IEEE Communications Letters* 2005;9(7): 643- 645.