RADIO PLANNING IN MULTIBEAM GEOSTATIONARY SATELLITE NETWORKS

C. Touati*, E. Altman*, J. Galtier*,‡, B. Fabre§, I. Buret§

* *INRIA BP93, 06902 Sophia Antipolis Cedex, France*

‡ *France Telecom R&D, 06921 Sophia-Antipolis Cedex, France*

§ *Alcatel Space Industries, Toulouse, France*

We consider the problem of how a geostationary satellite should assign bandwidth to several service providers (operators) so as to meet some minimum requirements, on one hand, and to perform the allocation in a fair way, on the other hand. In the paper, we firstly address practical issues (such as integrity constraints), whereafter we provide a computational method for obtaining an optimal fair allocation in polynomial time taking the practical issues into account.

keywords: bandwidth allocation, fairness, geostationary satellite.

1 Introduction

We consider a multi-spot geostationary satellite system for which a manager wants to assign bandwidth between various service providers (operators) that operate in different geographical areas. An actual assigned unit of bandwidth may correspond to different amounts of throughput, depending on many factors, and in particular on weather conditions. Indeed, during bad weather in an area, a local operator may have to use a larger part of its bandwidth for redundant information (a higher coding rate for error correction), and thus the effective throughput of information decreases. Therefore, if the objective was to maximize the global throughput, it would become non-profitable to assign bandwidth to operators in areas that suffer from bad weather if this bandwidth could be assigned to other operators instead. It is therefore interesting to understand and then to propose bandwidth assignment schemes that are more fair and do not systematically penalize operators that suffer from bad weather conditions.

The geographical area covered by a geostationary satellite is divided into hexagonal areas called *spots*. Each spot uses a certain frequency range denoted as its *color*. We denote $S(c)$ the set of spots of a given color c. In practice, as the number of available frequencies is limited, each satellite has a fixed number of colors that it can use. Two spots having the same color could interfere with each other. That is why they need to be geographically distant. The colors are statically assigned to the spots. A spot s is further divided into a set $Z(s)$ of *zones*. They are small enough to assume that the weather condition is the same in any point of a given zone. Then, every operator in a given zone uses the same coding rate and maximizing the global throughput does not penalize any operator with respect to any other within the same zone.

1.1 Technical framework

A central difficulty for solving such systems is that integrity constraints may arise: we might not be able to assign any value of bandwidth between the minimum and maximum given values. Instead, each operator in the set O of operators can be assigned one or more carriers in each zone among the set of N types of carriers: $T = \{1, \ldots, N\}$. Let B_t be the overall *bandwidth* of a type-t carrier. We assume that $B_1 > ... > B_N$.

B_t is not directly proportional to the actual throughput of information of a type-t carrier. Firstly, as mentioned before, the throughput depends on the coding rate, which may be different from one zone to another due to atmospheric conditions. Secondly, the effective throughput is lower due to overheads (around 10%) for signaling, frequency margins, etc; the percentage of overhead depends on the carrier type. To handle that, each carrier is associated to the utility $C_t(z, o)$. The utility is the value that operator o at zone z is willing to pay for a carrier of type t. It can be chosen as a function of the amount of redundancy (in the channel coding) which depends on the atmospheric conditions at each zone. Thus the utility can be made proportional to the actually assigned throughput, so the problem solved becomes how to fairly (or optimally) assign throughput.

We assume that there is a minimum and a maximum number of type-t carriers per zone z required by each operator o, denoted by $D_t^{min}(z, o)$ and $D_t^{max}(z, o)$ respectively. We assume that the minimum requirement can be satisfied; if not, our algorithm can be adapted to find new fair minima, see Appendix A.

The actual implementation of bandwidth assignment to users involves two phases. The first concerns the allocation of the global bandwidth to each operator in each zone. In that phase we ignore the problem of interferences between carriers of the same colors assigned to neighboring spots. In a second phase, the actual assigned bandwidth has to be translated to physical assigned resources such as slots, if a TDMA approach is used, or such as codes, if the CDMA approach is used. Here we do not address the second allocation problem, which involves complex combinatorial optimization (see, for instance[9]). Yet, in order to simplify that second part of the assignment, we study the possibility of adding the following *inter-spot compatibility condition* (ISCC). Roughly speaking, it can be viewed as (i) requiring a design of the same frequency plan to all spots of the same color, and then (ii) allow to replace a requirement of an operator for a given carrier j by assigning to it a carrier $t < j$ (i.e. with larger bandwidth) yet charging it for carrier j. If an operator cannot use in practice more bandwidth than it required, then this would mean wasting bandwidth.

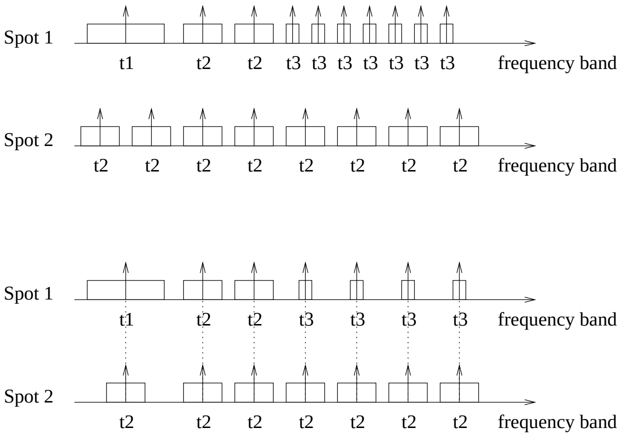

Fig. 1 Two allocations without and with ISCC.

An alternative way to formulate the ISCC is that a carrier location in the frequency plan is the same for all the spots of the same color. In other words, assume that $\mathcal{O}(f)$ is the set of operators that are assigned carriers that contain frequency f, and let t be the carrier with the largest bandwidth among $\mathcal{O}(f)$. Then the frequency ranges of all other $t' \in \mathcal{O}(f)$ must be contained in the frequency range of t, and moreover, an operator can only be assigned *a single carrier* in that range. We illustrate this in Fig. 1. On the upper part, spots 1 and 2 do not fulfill the ISCC. The lower part of the figure shows the two same spots with ISCC. We see that the carriers are then located at the same positions. But Spot 1 receives 4 carriers of type 3 instead of 7, and spot 2 receives 7 carriers of type 2 instead of 8. We shall show how to handle these constraints and what they cost with respect to fairness criteria.

1.2 Fairness and global optimization criteria

We are interested here in optimizing a fairness criteria (see[4–6]). There are several possible ways to understand and implement fairness concepts in our context, in the geographical sense. As we are interested in sharing bandwidth between operators, we may try to achieve:

- local fairness, i.e. in each zone,
- global fairness, i.e. in the total allocation,
- intermediate fairness, i.e. in a geographical area larger than a zone, such as a spot or a group of spots.

We clearly see that two spots having a different color are associated to two independent problems. In the following, we therefore compute the allocations independently among each color group.

Several concepts of fairness can be found in the literature. The most commonly found are :

The max-min fairness has been adopted by the ATM-forum[1] as the standard for throughput allocation for the best-effort traffic class ABR (Available Bite Rate). An allocation is max-min fair if and only if there is no possibility to allocate more excess throughput with respect to its minimum requirements to an operator unless we allocate less excess throughput to another operator that has less excess throughput than the previous one.

Proportional fairness characterizes TCP. Indeed, it has been shown that Vegas variant of TCP is proportionally fair.[7] A bandwidth assignment λ is proportionally fair if for any other assignment λ^*, the aggregate changes in the throughputs is zero or negative. That is to say, if X refers to the set of connections: $\sum_{x \in X} (\lambda_x^* - \lambda_x)/\lambda_x \leq 0$. It is shown in[8] that an assignment is proportionally fair if and only if it maximizes the product of the excess of bandwidth.

Global optimization. We may also consider the criterion of maximizing the total used throughput although we saw that it would lead to an unfair solution. It is called *Global Optimization*.

All these criteria produce Pareto optimal solutions ([3]). An allocation λ is said to be Pareto optimal if there is no other allocation that strictly dominates it. This means that there is no allocation where all operators get an allocation greater than or equal to λ with at least one operator getting strictly more than according to λ. The Pareto definition assures that no bandwidth is wasted.

A general fairness criterion that covers all the previous one is due to.[8] Given a positive constant $\alpha \neq 1$, and a color c, consider the maximization of:

$$\frac{1}{1-\alpha} \sum_{o \in O} \left[\sum_{\substack{s \in S(c), z \in Z(s) \\ t \in T}} (D_t(z,o) - D_t^{min}(z,o)) \right]^{1-\alpha}$$

subject to the problem's constraints. Since the utility function is concave and the constraints are linear, this defines the unique α-throughput allocation. It corresponds to the *globally optimal allocation* when $\alpha = 0$, the *proportional fairness* when $\alpha \to 1$, the *harmonic mean fairness* when $\alpha = 2$, and the *max-min* allocation when $\alpha \to \infty$ (see[8]).

Similarly, we can define the weighted α-throughput allocation as the unique solution of maximizing:

$$\frac{1}{1-\alpha} \sum_{o \in O} \left[\sum_{\substack{s \in S(c),\\ z \in Z(s)\\ t \in T}} C_t(z,o)(D_t(z,o) - D_t^{min}(z,o)) \right]^{1-\alpha} \tag{1}$$

We notice by computing the limit of this criterion when $\alpha \to 1$, that there is no "weighted proportional fairness" in the sense of equation 1. This is a remarkable property of proportional fairness: multiplying the value of a connection by a given factor does not modify the allocation (see[10]).

1.3 Problem formulation

We focus on the fairness criterion obtained as solutions of maximizing:

$$\sum_{\substack{t \in T, s \in S(c),\\ z \in Z(s), o \in O}} \frac{\left[C_t(z,o)(D_t(z,o) - D_t^{min}(z,o))\right]^{1-\alpha}}{1-\alpha} \tag{2}$$

This agrees with the previous definition in the case of global optimization. It also agrees with the previous one if we identify a sub-operator to be responsible for the demand of a given type of a given operator at a given zone, and then apply the previous definitions of fairness to the sub-operators. Thus this assignment is Pareto optimal and has essential fairness properties.

Finally, for a given color c, our problem is to compute the optimal J_t, $t = 1, \ldots, N$ and $D_t(z,o)$ where J_t is the number of carriers of type t in a spot of color c and $D_t(z,o)$ is the number of type-t carriers assigned on zone z to operator o.

The constraints of our problem can be summarized as follows :

$$\begin{array}{ll} \forall z \in Z(c), \forall o \in O, & \\ \forall t \in T, & D_t^{min}(z,o) \le D_t(z,o) \le D_t^{max}(z,o), \\ \forall s \in S(c), \forall t \in T, & \displaystyle\sum_{i \in T s.t. i \le t} J_i \ge \sum_{z \in Z(s), o \in O} D_t(z,o), ^* \\ \forall t \in T, & J_t \ge 0, \\ & \displaystyle\sum_{i \in T} J_i B_i \le B. \end{array}$$

2 Our allocation algorithm

This part focuses on solving the ISCC constraint. We firstly assume that the number of carriers of each type is known and fixed for all spots of a given color, and determine to what operator we should assign them. Then, we show how to determine the number of carriers of each type. Finally, we show how these two components are combined to obtain our optimal fair allocation.

2.1 Assignment of the carriers to operators in a single spot

Within a spot s, let p be the total number of demands, namely, $p = \sum_{z \in Z(s), o \in O, t \in T} D_t^{max}(z,o)$. We consider the set of these demands as a single dimension array (of size p) and note their types $t_1, \ldots, t_p$ (where $t_i \in \{1, ..., N\}$), and associate with them the values $v_1, \ldots, v_p$. We will see in Section 2.3 that an appropriate choice of values v_i leed us to the desired fairness criterion. The demands are sorted so that $v_1 \ge \cdots \ge v_p$. Our objective is to maximize the total utility.

We work with a fixed given vector $J = (J_1, ..., J_N) \in \mathbf{N}^N$. J_i corresponds to the number of carriers of type i that can be assigned to demands of type j where $j \ge i$.

We want to find *Accept*, a subset of $\{1, \ldots, p\}$, that satisfies the constraint of available carriers (i.e. the number of carriers of type i given by J_i, $i = 1, ..., N$) and maximizes the value $V = \sum_{i \in Accept} v_i$.

The algorithm presented in Fig. 2 solves this problem. The idea is as follows: we add one by one carriers, starting by the most profitable ones. k corresponds to the last level of carrier type that has been filled. More precisely, when all available carriers of type $1, 2, ..., j$ have been used (the number of assigned carriers of these types has reached $J_1 + ... + J_j$) then $k = j$. For each new carrier we check whether its type is larger than k. In this case it is accepted and we update the value of k. Otherwise it is rejected.

This algorithm gives the assignment that maximizes the sum V with respect to the vector $J = (J_1, ..., J_N) \in \mathbf{N}^N$ in a given spot. For a proof of its optimality, see Appendix B.

$k := min\{i$ such that $J_i > 0\} - 1$
for i:=1 to p
  begin
    if $t_i > k$ then
      begin
        accept[i]:=1
        for l:=$k+1$ to N
          if $\sum_{m \le i \text{ and } t_m \le l}$ accept[m] $= \sum_{n \le l} J_n$
            then $k := l$;
      end
      else accept[i]:=0;
  end;

Fig. 2 An assignment algorithm to maximize the sum of the values.

*This results from the fact that we can assign to an operator a carrier of greater type than it requires.

The algorithm produces a vector "accept" of p boolean elements, and the set Accept consists of carriers whose corresponding "accept" value is 1.

2.2 Obtaining a global J

In order to determine the most profitable global $J = (J_1, ..., J_N) \in \mathbf{N}^N$, we now consider the following global assignment procedure in a given color.

For each admissible J (satisfying the minimum demand), we find in each spot the best assignment of carriers to operators using the algorithm of Section 2.1. We evaluate the quality of the assignment allowed by J using our global fairness criteria. The J having the best value is selected.

2.3 Basic steps for the general solution

Therefore, the solution of our problem can be found in three steps:

Satisfaction of the minimum requirements. For each spot s we set

$$J_t^s = \sum_{o \in O} \sum_{z \in Z(s)} D_t^{min}(z, o),$$

and set $J^* = max(\{J^s\}_{s \text{ spot}})$.[†] If $J^* \in \mathcal{B}$ then all the minimum requirements can be satisfied. Otherwise, the system is non-feasable. From a practical point of view, this step can be done by using the algorithm described in Section 2.1 and setting $v_i = 1$ for the minimum demands, and 0 for the other ones.

Maximization of the surplus beneficiaries. Equation 2 shows that, for $\alpha \geq 1$, if for some t, z, o we get $D_t(z, o) = D_t^{min}(z, o)$, then the optimization criteria given by Eq. 2 goes to $-\infty$. Therefore, we aim at minimizing its degree of nullity by maximizing the number of operators to whom we can assign strictly more than their required minimum. For this, we use the algorithm of Section 2.1, and set the value v of a demand i to:

$$\begin{cases} 2 & \text{for a minimum (i.e. a required) demand} \\ 1 & \text{for the smallest demand above the miminum} \\ & \text{requirement of each beneficiary}^{\ddagger} \\ 0 & \text{otherwise} \end{cases}$$

We then restrict our optimization problem only to those operators.

Fair redistribution of the values. We use the algorithm of Section 2.1. We set the value v_i of any minimum demand i to ∞ (or any upper bound). If $\alpha \geq 1$, we set the value of the smallest demand to ∞ for the operators to whom we can assign more than their minimum requirements. We then fix the value v_j of each non-assigned demand j ($D_t^{min}+2 \leq j \leq D_t^{max}$) to:

$$\begin{cases} \log\left(\frac{j - D_t^{min}(s,o)}{j-1-D_t^{min}(s,o)}\right) & \text{if } \alpha = 1^{\S} \\ C_t(s,o)^{1-\alpha}\left[(j - D_t^{min}(s,o))^{1-\alpha}\right. & \\ \quad \left. -(j-1-D_t^{min}(s,o))^{1-\alpha}\right] & \text{otherwise.} \end{cases}$$

If $\alpha < 1$, we use the same equations for all demands above D_t^{min} ($D_t^{min} + 1 \leq j \leq D_t^{max}$)

3 Numerical Figures

3.1 The considered network

We consider an example with 32 spots and 4 colors, so that 8 spots receive the same color.

Each spot is covered by one MF-TDMA channel (whose bandwidth is 36MHz). It can be used with five types of carriers, taking the bandwidths of 6000, 3000, 1500, 750, and 187.5 in kHz. The spacing is of 50%, therefore the bandwidth would allow a maximum number of 4 carriers having the first type, 8 of the second, 16 of the third, 32 of the fourth and 128 of the last one - if all the bandwidth was used for only one type of carrier. Each spot is divided into three zones, and the value C_t of the carrier type in one of the spots is given by:

$C_t(z, o)$	Zone z=0	Zone z=1	Zone z=2
Type t=1	4920.000000	4373.328613	3280.000000
Type t=2	2460.000000	2186.664307	1640.000000
Type t=3	1230.000000	1093.332153	820.000000
Type t=4	615.000000	546.666077	410.000000
Type t=5	153.750000	136.666519	102.500000

These different values of C_t can represent various estimations of the types of carriers by the operators. In particular, radio signal fading can require an increasing of the Viterbi coding from 3/4 to 1/2, which changes the final rate and the service provided by the operator.

We suppose that there are 40 operators and that each may use at most 3 spots. There may be at most two operators per spot, and any operator in a spot may have 1, 2 or 3 zones. The same zone may be assigned to more than one operator. To achieve this, we choose three different spots among 32 for each operator. If one of these spots already has two operators, we

[†] We can consider a partial order among the assignments: we say that an assignment J' is greater than another assignment J (and we write $J' \succeq J$) if and only if

$$\begin{aligned} J_1' &\geq J_1, \\ J_2' + J_1' &\geq J_2 + J_1, \\ \dots & \quad \dots \\ \sum_{t=1}^{N} J_t' &\geq \sum_{t=1}^{N} J_t. \end{aligned}$$

We have the following property (the proof can be found in Appendix C): given two assignments J and J', there is a unique assignment $J^* = \max(J, J')$ such that $J^* \succeq J$, $J^* \succeq J'$, and for any assignment J'' such that $J'' \succeq J$ and $J'' \succeq J'$, we have $J'' \succeq J^*$.

[‡] which can be an operator per zone or an operator per spot or an operator per zone and per type of carrier, etc.

[§] One can notice that the value v. is independent of C. This is because there is no "weighted proportional fairness", see Section 1.2.

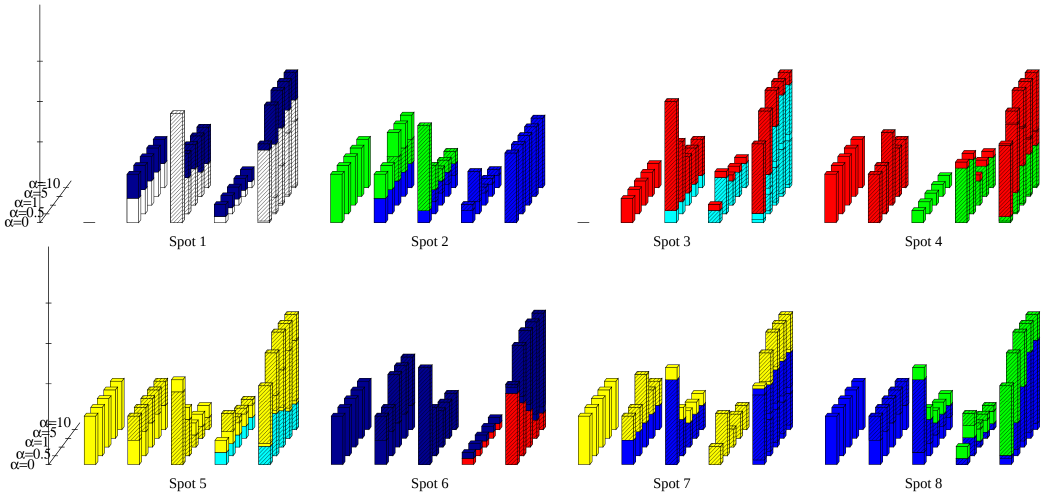

Fig. 3 Fair allocation of bandwidth with ISCC

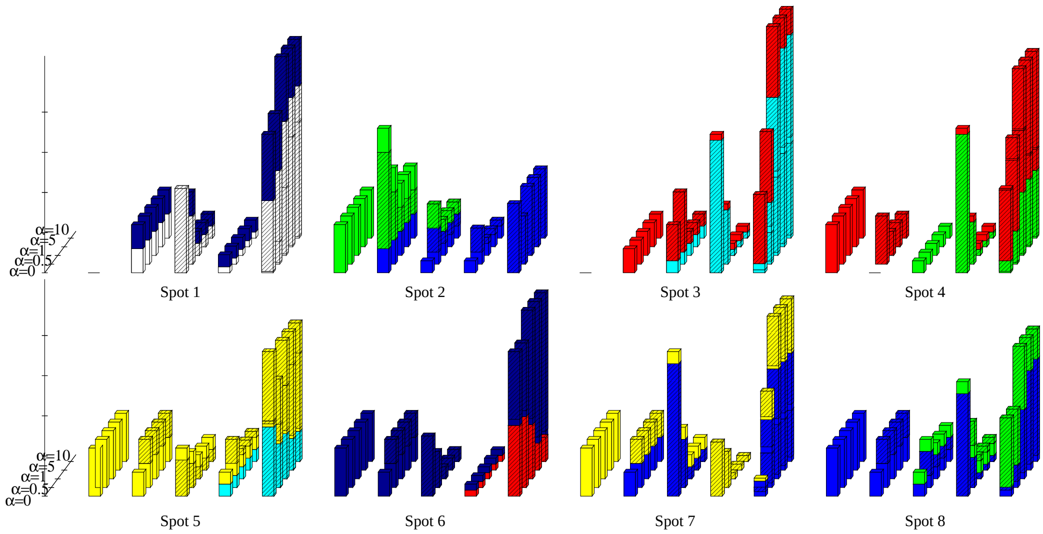

Fig. 4 Fair allocation of bandwidth without ISCC

simply cancel the request of the new operator. Within each selected spot, the operator chooses at random if it will use each type of carrier with probability 1/2. If it uses no carrier in the end, its request is withdrawn for this spot. We set the maximum number of carriers to infinity. The minimum number of carriers is chosen at random to occupy between 20% and 50% of the bandwidth divided by the amount of types of carriers used by this operator. We assumed in the numerical problem that the demands for carriers by an operator o in a zone z is given by three objects: a single carrier type t, a maximum $d_{\max}$ and a minimum $d_{\min}$ number of required carriers. These objects are assumed to have the following interpretation: the operator requires *exactly* $D_t^{\min}(z,o) = D_t^{\max}(z,o)$ carriers of type t and wishes to have any number not greater than $d_{\max} - d_{\min}$ of carriers of type $\min(t+1, N)$. In other words, $D^{\min}_{\min(t+1,N)} = 0$, $D^{\max}_{\min(t+1,N)} = d_{\max} - d_{\min}$, and for all $s > \min(t+1, N)$, $D_s^{\min} = D_s^{\max} = 0$. E.g, if an operator asks for a minimum of 10 carriers of type-1, and a maximum of 20 carriers, it will receive the type-1 carriers it requested (that is 10) and will receive in addition at most 10 more carriers of type-2.

3.2 Numerical results

Results are given by Figs 2.3 and 2.3. We plot the bandwidth given to each operator in the different spots having the same color, and for all types of carriers. Each operator receives a distinct color. In each spot we accumulate the bandwidth assigned to the operators for a particular spot and carrier. The height of the bar is proportional to the bandwidth. Hence, an additional 2-type carrier will increase more the bar than an additional 3-type carrier. We plot the minimum demand plus the surplus, the surplus being on a separate hachured zone, and as said earlier, on a decreased type.

Figs 2.3 and 2.3 show clearly that, as α grows,

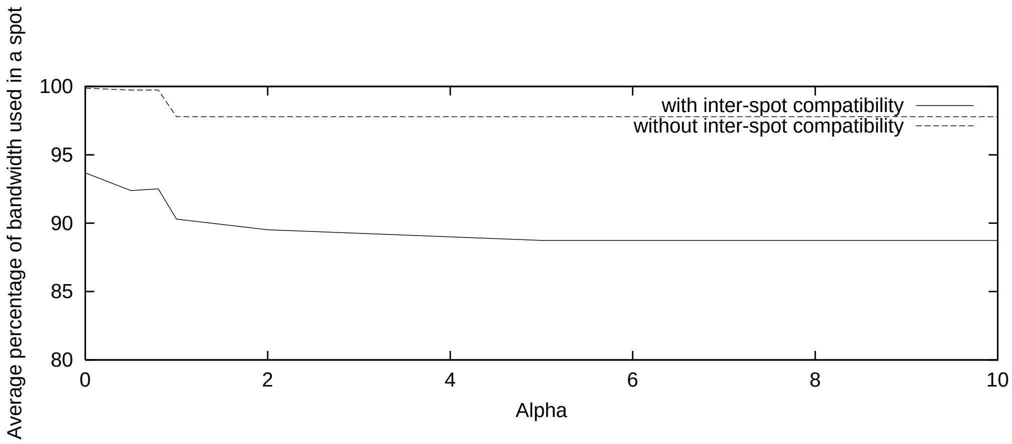

Fig. 5 Bandwidth efficiency with ISCC

carriers of higher indexes are favored compared to carriers of lower indexes. Indeed, the fairness criterion redirects progressively the global optimization policy ($\alpha = 0$) to a max-min policy ($\alpha \to \infty$). Also, as α grows, the total throughput is reduced which is the natural price of the fairness policy. An interesting remark also can be draw from Fig. 5. We note that ISCC constraints lead to more waste of the bandwidth as α grows, up to 10% but no more. This comforts us in the fact that the cost of the ISCC remains reasonnable.

Conclusion

In this paper we have studied different optimality and fairness criteria that can be treated uniformly using a single optimization problem with a parameter α that corresponds to the type of fairness or optimality we wish to achieve. All concepts lead to Pareto optimal solutions. We proposed an efficient computational method to achieve the solution in a polynomial time, taking into account ISCC. The method allows to obtain a large range of solutions depending on the fairness parameter. We have also shown that the ISCC constraint presents an acceptable cost to the bandwidth, which can lead to new perspectives of resource allocation algorithms.

References

[1]The ATM forum Technical Committee, *Traffic Management Specification,* Version 4.0, April 1996.

[2]D. Bertsekas and R. Gallager, *Data Networks*, Prentice-Hall, 1987.

[3]T. Bonald and L. Massouillé, “Impact of fairness on internet performance”, *ACM Sigmetrics*, 2001.

[4]J. Crowcroft and P. Oechslin, “Differentiated end-to-end Internet services using a weighted proportional fair sharing TCP”, *Computer Communications Review*, 28(3):53–67, July 1998.

[5]F. P. Kelly, A. Maulloo and D. Tan, “Rate control for communication networks: shadow prices, proportional fairness and stability”, *Journal of the Operational Research Society* **49**, 1998.

[6]L. Massouillé and J. W. Roberts, “Bandwidth sharing and admission control for elastic traffic”, *Telecommunication Systems*, 2000.

[7]S. Low, L. Peterson and L. Wang, “Understanding TCP Vegas: a duality model”, *ACM SIGMETRICS*, pp 226-235, June 2001.

[8]J. Mo and J. Walrand, ”Fair End-to-End Window-based Congestion Control”. SPIE ’98, International Symposium on Voice, Video and Data Communications, 1998.

[9]T. Toker, E. Altman, J. Galtier, C. Touati, I. Buret, B. Fabre, C. Guiraud, ”Slot Allocation in a TDMA Satellite System: Simulated Annealing Approach ”, *AIAA International Communication Satellite System Coference*, April 2003.

[10]C. Touati, E. Altman and J. Galtier, ”On fairness in bandwidth allocation”, *INRIA Research Report RR-4269*, September 2001.

Appendix

A Definitions and computation of fairness for allocation of minimum throughput

In this phase we wish to see whether all minimum required throughput can be allocated.

This can be done by a linear program in which we maximize the total throughputs allocated to all operators, adding the upper constraint that no operator in any zone should get more than its minimum requested throughput. If we denote by λ_s^{min} the minimum throughput guaranteed to operator s and define λ_s the allocated throughput to that operator, then this constraint is given by $\lambda_s \leq \lambda_s^{min}$ for all $s \in S$.

We note that the linear program may end up assigning to some operators their required minimum throughput, whereas others may receive much less than their required minimum. We may consider this as an unfair solution, and prefer a solution in which some more operators get less than their required minimum throughput but the less satisfied operators end up closer to their requirement than in the previously mentioned allocation. In other words, we may try to obtain new values of the minima in a "fair" way. We then may proceed as follows using one of the following two approaches.

1st approach: maximizing throughputs We could use the *same definitions of fairness as before*, set the minimum rate for all operators to zero and the maximum throughput for operator s to $\lambda_s^{\min}$. Then the computations are the same as before.

2nd approach: minimizing the undelivered bandwidth We could use the *same definitions of fairness as before*, yet instead of applying them to the excess capacity (which we wished to maximize), we can apply them to the amount of bandwidth not delivered $\tilde{\lambda}_s := \lambda_s - \lambda_s^{min}$. We note that this is a non-positive quantity (that we still wish to maximize), and thus we cannot apply anymore some of the formulas for computing the fairness criteria (see, in particular, the one for the proportional fairness).

To avoid that problem we may simply consider the variables $|\tilde{\lambda}_s| = -\tilde{\lambda}_s$ which are positive, and then redefine the fairness concepts accordingly. The max-min fairness criterion will transfer to a min-max criterion in the obvious manner. For the proportional fairness, we can still use the definition with $\tilde{\lambda}_s$ replacing λ_s, or equivalently work with $|\tilde{\lambda}_s|$ and define $|\tilde{\lambda}|$ to be proportionally fair if it is feasible (satisfies the constraints) and if for any other feasible assignment $\tilde{\lambda}^*$, the aggregate of proportional changes is zero or positive

$$\sum_{s \in S} \frac{\tilde{\lambda}_s^* - \tilde{\lambda}_s}{\tilde{\lambda}_s} \geq 0. \tag{3}$$

One can then show that the proportional fairness has similar properties as before: it is an assignment that *minimizes* the quantity $\prod_s |\tilde{\lambda}_s|$. Equivalently, it is an assignment that minimizes $\sum_s \log |\tilde{\lambda}_s|$. Unfortunately this problem is computationally more complex than the maximization of these quantities, since unlike the case of maximization, a local minimum may not be a global minimum. Moreover, the uniqness of the solution is not guaranteed anymore.

A similar complication occurs when redefining the general α-fairness: the maximization has to be changed to minimization and we are not guaranteed to get a unique solution.

We conclude that the first approach is preferable, as it is much easier to solve and it guarantees a unique solution.

B Optimality of the algorithm

The aim of this subsection is to show that our algorithm finds the optimal assignment of carriers for a fixed given vector J.

In the general context, we consider in the following a finite set I of positive valued elements. Let v_i denote the value of element i.

Definition B.1 *Let A and B be two disjoint subsets of I. Then define $Union(A, B, k)$ as a subset H of at most k elements of $A \cup B$ such that $\sum_{i \in H} v_i$ is maximum.*

Note that if $I = \{1, \ldots, p\}$ and $v_1 \geq \cdots \geq v_p$, H can be obtained by the k smallest indices of $A \cup B$.

Definition B.2 *Let A and B be two subsets of I. We say that A dominates B iff $|A| \geq |B|$ and there exists an injective mapping φ from B to A such that*

$$\forall i \in B \qquad v_i \leq v_{\varphi(i)}.$$

Proposition B.1 *Let A, B, C and D be three subsets of I such that*

- *A dominates B,*
- *C dominates D,*
- *$A \cup B$ and $C \cup D$ are disjoint.*

Then, for all $k \in \mathbb{N}$, $Union(A, C, k)$ dominates $Union(B, D, k)$.

Proof. Let φ_1 be a dominating maping from B to A, and φ_2 one from D to C. Define ψ from $B \cup D$ to $A \cup C$ as follows:

$$\begin{array}{lll} \psi(v) & \mapsto & \varphi_1(v) \quad \text{if } v \in B, \\ \psi(v) & \mapsto & \varphi_2(v) \quad \text{if } v \in D. \end{array}$$

Obviously the restriction of ψ to $Union(B, D, k)$ generates a set of at most k elements in $A \cup C$, that is by definition dominated by $Union(A, C, k)$. □

Now consider that each element i of I is given a type t_i and denote

$$A_p = \{i \in A: \ t_i = p\}.$$

Obviously our allocation problem consists of finding a set A such that

$$\forall k \in \{1, \ldots, N\} \qquad |A_1 \cup \cdots \cup A_k| \leq J_1 + \cdots + J_k.$$

Theorem B.1 *Let $A^{(k)}$ be constructed as follows:*

$$\begin{aligned} A^{(0)} &= \emptyset \\ A^{(k+1)} &= Union(A^{(k)}, I_{k+1}, J_1 + \cdots + J_{k+1}) \end{aligned}$$

then $A^{(N)}$ achieves a solution of our allocation problem that dominates all the other solutions.

Proof. Obviously $A^{(N)}$ is a solution of our allocation problem. Let B be an alternative solution. We shall proove by induction that $A^{(k)}$ dominates $B_1 \cup \cdots \cup B_k$.

In fact, we have $B_1 \cup \cdots \cup B_{k+1} = Union(B_1 \cup \cdots \cup B_k, B_{k+1}, J_1 + \cdots + J_{k+1})$. Clearly I_{k+1} dominates B_{k+1}, and by induction $A^{(k)}$ dominates $B_1 \cup \cdots \cup B_k$. Hence the result by proposition B.1. □

C Order among the assignments

We say that an assignment J' is greater than another assignment J (and we write $J' \succeq J$) if and only if

$$\begin{aligned} J'_1 &\geq J_1, \\ J'_2 + J'_1 &\geq J_2 + J_1, \\ &\vdots \\ \sum_{t=1}^{N} J'_t &\geq \sum_{t=1}^{N} J_t. \end{aligned}$$

Obviously the relation $\succeq$ is a partial order.

We suppose here that an operator has a demand in some spot of a vector of j carriers which has to be satisfied. We assume that by giving more, the operator would be satisfied as well. Hence a demand for j t-carriers is considered to be satisfied when assigning instead j k-carriers, with $k \leq t$.

The advantage of offering more throughput to an operator is illustrated in the following example:

Example C.1 *Spot 1 and 2 have the same color.*
Operator 1 pays 30\$ for a carrier of type 1 in spot 1.
Operator 2 pays 80\$ for a carrier of type 2 in spot 1.
Operator 3 pays 50\$ for a carrier of type 2 in spot 2.
Operator 4 pays 20\$ for a carrier of type 2 in spot 2.

We can satisfy all these demands by a carrier of type 1 and a carrier of type 2 in both spots.

The above "upgrade" scheme implies that a demand J of an operator in a zone, can be satisfied with an assignment J' satisfying $J' \succeq J$.

We have the following key property of the above ordering:

Proposition C.1 *Given two assignments J and J', there is a unique assignment $J^* = \max(J, J')$ such that $J^* \succeq J$, $J^* \succeq J'$, and for any assignment J'' such that $J'' \succeq J$ and $J'' \succeq J'$, we have $J'' \succeq J^*$.*

Proof We set

$$J^*_1 = \max\{J_1, J'_1\}$$
$$\vdots$$
$$J^*_i = \max\left\{\sum_{k=1}^{k=i} J_k, \sum_{k=1}^{k=i} J'_k\right\} - \sum_{k=1}^{k=i-1} J^*_k$$
$$\vdots$$
$$J^*_N = \max\left\{\sum_{k=1}^{k=N} J_k, \sum_{k=1}^{k=N} J'_k\right\} - \sum_{k=1}^{k=N-1} J^*_k$$

which satisfies the required properties. ∎

The interpretation of the above proposition in our context is as follows. Given minimum rate constraints J and J' for two operators, any assignment J'', that is candidate for being a (feasible) fair assignment, should satisfy $J'' \succeq J^*$. This extends in an obvious way to any number (larger than two) of operators. The value J^*, which can be determined explicitly as done in the proof of Proposition C.1, can now serve as the starting point for fair assignments algorithms: we first determine the unique element J^* and then find the best assignment J'' s.t. $J'' \succeq J^*$.

Unfortunately, the bandwidth assigned to our service is limited. There is therefore an additional constraint on the maximum number of carrier of each type. The constraint is of the form:

$$\sum_{i=1}^{N} B_i J_i \leq C \tag{4}$$

with C the total bandwith and B_i the bandwith of a i-carrier. We denote by $\mathcal{B}$ the set of acceptable assignments with respect to the bandwidth. Obviously if $J \succeq J'$ and $J \in \mathcal{B}$ then $J' \in \mathcal{B}$. Nevertheless, it is possible that for some $J \in \mathcal{B}$ and $J' \in \mathcal{B}$ we have $\max(J, J') \notin \mathcal{B}$. In other words, an assignment can be locally but not globally feasible.