Exploiting a Common Property Resource under a Fairness Constraint: a Case Study

Michel Lemaitre and Gerard Verfaillie
ONERA CERT
2, avenue Edouard Belin (BP 4025)
31055 - Toulouse Cedex - Prance

Nicolas Bataille
Centre National d'Etudes Spatiales
18, avenue Edouard Belin
31055 Toulouse Cedex Fiance

Abstract

Resources co-funded by several agents must be exploited in such a way that three kinds of constraints are met: (1) physical problem (hard) constraints; (2) efficiency constraints, aiming at maximizing the satisfaction of each agent; (3) a fairness constraint, which is ideally satisfied when each agent receives an amount of the resource exactly proportional to its financial contribution. This paper investigates a decision problem for which the common property resource is an earth observation satellite. The problem is to decide on the daily selection of a subset of pictures, among a set of candidate pictures which could be taken the next day considering the satellite trajectory. This subset must satisfy the three kinds of constraints stated above. Although fair division problems have received considerable attention for a long time, especially from microeconomists, this specific problem does not fall entirely within a classical approach. This is because the candidate pictures may be incompatible, and because a picture is only of value to the agent requesting it. As in the general case, efficiency and fairness constraints are antagonistic. We propose three ways for solving this share problem. The first one gives priority to fairness, the second one to efficiency, and the third one computes a set of compromises.

1 Introduction

Due to their cost, large research or industrial projects are often co-funded by several agents (countries, companies, entities ...). Space projects such as earth observation satellites, space stations or space probes are good examples. Once constructed and made operational, the common property resource must be exploited and shared in a way which satisfies three kinds of constraints :

- *physical* constraints: the exploitation of the resource must obey hard constraints;
- *efficiency* constraints: each agent wants to get the highest possible satisfaction in return;
- a *fairness* constraint: each agent must get a return on investments proportional to its financial contribution to the project; the better the proportionality of returns is achieved, the more the share quality improves.

The first kind of constraints must absolutely be met (hard constraints) whereas the two others are preference constraints (soft constraints). As it can be easily guessed, the efficiency and fairness constraints are antagonistic: the search for a perfect share may lead to poorly efficient decisions, and conversely, decisions which maximize the global satisfaction of agents are often unfair. So, a compromise between the best satisfaction of both constraints must be found.

The usual case involving only one agent (in which case there is no share problem) is a difficult combinatorial discrete optimization problem (NP-hard). Nevertheless, it is a perfectly well stated problem. The multiagent case is also a discrete combinatorial problem, but is actually a multi-objective optimization problem [Keeney and Raiffa, 1976]; the first difficulty arises when searching for a meaningful and principled definition of a good compromise between efficiency and fairness.

This article sums up a study, the aim of which was to propose methods to solve a specific share problem, namely the fair and efficient exploitation of an earth observation satellite owned in common by several agents. It is organized as follows. The next section sets the problem more formally. Then we present three quite different methods devoted to the resolution of this share problem. These methods have been simulated on the basis of the expected data for the future Spot5 satellite. The section after reports these simulations. Lastly, we state our conclusions.

2 An Earth Observation Satellite Scheduling and Sharing Problem

The studied problem is the following: an earth observation satellite, co-funded by several agents, is exploited in common. These agents make daily requests for pictures they would like to be taken by the satellite. Roughly speaking, the problem consists in selecting each day, among the set of candidate pictures which could be taken the next day considering the satellite trajectory, a subset of pictures which satisfies all the physical constraints, maximizes the satisfaction of the agents, and respects as far as possible a fairness constraint. Such a selection will be called a *decision.* The satisfaction and fairness constraints will be taken into account over a fixed interval of several days.

Let us describe more formally the problem. First, the data:

- there are n agents; in our real-world problem, n is typically in the range 3 to 8;
- D_{ik} is the set of pictures requested by the agent i for the day k; let

$$D_{\bullet k} \stackrel{\text{def}}{=} \bigcup_{i=1}^{n} D_{ik}; \qquad (1)$$

  the size of a $D_{\bullet k}$ is averaging 200;
- the sets D_{ik} are pairwise disjoint : $D_{ik} \bigcap D_{jk} = \emptyset$ for all i, j;
- each picture in $D_{\bullet k}$ could be taken the day k, but all pictures cannot be taken because there are incompatibilities between them: some physical hard constraints must be met (for example no more than m pictures can be taken at once, provided there are only m instruments on board; a transition time between two pictures taken by the same instrument must be respected; on board memory is limited ...); a subset $X \subseteq D_{\bullet k}$ is said *admissible* if all pictures in X satisfy hard constraints (pictures are compatible) and hence can all be taken the considered day;
- $w(x)$ is the weight of the picture x; it is freely set by the agent requesting the picture, and reflects its importance for the agent;
- $q = (q_1, \cdots, q_n)$, with $\sum_{i=1}^{n} q_i = 1$ is the quota vector: q_i is proportional to the financial investment of the agent i.

We characterize now the decisions that we are looking for. Each day $k-1$, the demands D_{ik}, with corresponding weights, are collected and we must compute the sets of pictures A_{ik} which will be shot for the agent i the day k. These A_{ik} are such that:

- $A_{ik} \subseteq D_{ik}$ (note that A_{ik} are disjoint);
- let

$$A_{\bullet k} \stackrel{\text{def}}{=} \bigcup_{i=1}^{n} A_{ik}; \qquad (2)$$

  $A_{\bullet k}$ must be admissible;
- the *cumulative satisfaction* of each agent, measured over a given interval of days I ending on the day k must be as high as possible (efficiency constraints); the satisfaction of the agent i the day k is measured by the quantity $s(A_{ik})$ where

$$s(X) \stackrel{\text{def}}{=} \sum_{x \in X} w(x); \qquad (3)$$

  hence, the cumulative satisfaction over I for the agent i is

$$cs_i \stackrel{\text{def}}{=} \sum_{k \in I} s(A_{ik}); \qquad (4)$$

  These satisfactions need to be normalized over agents, if we compare or aggregate them.
- the "quality of the share" over I (to be formalized later) must be as high as possible (fairness constraint).

n	number of agents.
i	agent index."
k	day index.
I	interval of days on which satisfactions and costs are taken into account,
D_{ik}	pictures requested by agent i day k.
A_{ik}	pictures obtained by agent i day k.
$\mathbf{q}$	quota vector, $q = (q_1, q_2, \ldots, q_n)$.
$w(x)$	weight of picture x.
$s(X)$	satisfaction for an agent receiving the set X of pictures. See eq. 3.
cs_i	cumulative satisfaction. See eq. 4.
cs_i^M	cumulative maximal satisfaction. See eq. 7.
$c(x)$	cost of picture x.
cc_i	cumulative cost. See eq. 16.
$\mathbf{cc}$	vector of cumulative costs. See eq. 17
j	quality of share criterion. See eq. 21
gcs	global cumulative satisfaction. See eq. 22

Table 1: Main symbols used in this paper.

The problem above is stated as a sequence of multi-objective optimization problem instances. However, the fairness constraint is not yet formally stated. We have investigated three quite different methods devoted to the resolution of this share problem (that is general schemes for computing the A_{ik}). Each one is based on a particular way of taking into account the fairness constraint and the necessary compromise with the efficiency constraints. The first two methods reduce the problem to a sequence of mono-objective optimization problem instances, whereas the third one keeps the multi-objective aspect.

3 Fairness first

The first method searches for fairness first, and then for efficiency. The entitlement to use the resource is shared by allocating observation windows to each agent in turn. Observation windows are merely sequences of successive orbits of the satellite. Each day, the agent i is given the right to freely exploit about $Q_i \cdot N$ orbits, where N is the number of orbits daily covered by the satellite. Observation windows can be assigned to agents on the basis of a fixed repetitive procedure. This procedure and the trajectory of the satellite are such that each agent gets opportunity to shoot any place in the world within a bounded number of days.

Following this method, the whole problem can be cast into a set of optimization problem instances, one for each agent each day, because each agent knows in advance his time windows. Assuming that each $x \in D_{ik}$ belongs to the window assigned to agent i the day k, the successive optimization problem instances consist in maximizing the satisfaction of agents by finding

$$A_{ik} = \text{argmax}\{s(X) | X \subseteq D_{ik}, X \text{admissible}\}. \quad (5)$$

This problem can be seen as a combination of discrete *constraint* and *optimization* problems. General frameworks such as the *Semiring* and *Valued Constraint Satisfaction Problems* frameworks [Bistarelli *et al.*, 1995; Schiex *et al.*, 1995] have been recently designed to capture such mixed problems. Powerful complete and incomplete algorithms, associated to these frameworks, are now available, and research in this area is very active (Freuder and Wallace, 1992; Wallace, 1994; Verfaillie *et al.*, 1996; Larrosa *et al.*, 1998].

Our simulations are based on the Valued CSP framework. Almost all windows can be solved to optimality, using a sophisticated algorithm. These simulations show for this method a very good quality share: the number of pictures effectively selected and assigned to each agent is very close to a number proportional to its quota. But the decisions are clearly inefficient, when compared with those resulting from the two following methods, as reported in section 6.

4 Efficiency first

The second method considers the opposite view: first efficiency, fairness if possible. It is based on three main ideas:

1. for efficiency, maximize each day a linear combination of individual satisfactions of the agents;
2. for fairness, choose this combination in a way favoring the fairness constraint;
3. check that each agent has obtained a *fair share.*

The last point is borrowed from the literature on fair division [Young, 1994; Moulin, 1995; Brams and Taylor, 1996]: in this method, we postulate that a decision is fair when each agent receives at least a minimal fair share, defined for the agent i as q_i times the satisfaction it would get if it were the only user of the resource. More formally, the fairness constraint is considered to be satisfied if

$$cs_i \geq q_i \cdot cs_i^{\text{M}}, \ i = 1, \ldots, n, \quad (6)$$

$$\text{with } cs_i^{\text{M}} \stackrel{\text{def}}{=} \sum_{k \in I} s^{\text{M}}(D_{ik}), \quad (7)$$

$$s^{\text{M}}(D_{ik}) \stackrel{\text{def}}{=} \max\{s(X) | X \subseteq D_{ik}, X \text{admissible}\}. \quad (8)$$

We now turn to the determination of the linear combination of individual satisfactions to be maximized (points 1 and 2). For the moment, assume that the quotas are equal (all agents have equal rights over the common resource). The clue is to consider the weights of pictures as monetary bids. As a first approach, we could select pictures in such a way that the sum of bids for selected pictures, namely

$$s(A_{\bullet k}) = \sum_{x \in A_{\bullet k}} w(x) = \sum_{i=1}^{n} s(A_{ik}) \quad (9)$$

is maximum (under admissibility constraints). In this way, the higher the bid for a picture is, the more this picture gets some chance of being selected. But recall that weights are freely fixed by agents. The above function to be maximized must be corrected, in order to take this fact into account. In other words, we must make satisfactions comparable by *normalizing* them. So instead, the function to be maximized will be:

$$s'(A_{\bullet k}) \stackrel{\text{def}}{=} \sum_{i=1}^{n} s'(A_{ik}) \stackrel{\text{def}}{=} \sum_{i=1}^{n} \alpha_{ik} \cdot s(A_{ik}), \quad (10)$$

where the coefficients α_{ik} have to be determined. The principle of the normalization is the following : *the maximum of normalized individual satisfaction that an agent would get if it were the only user of the resource is equal for all agents.* Formally stated :

$$s'^{\text{M}}(D_{ik}) = 1, \forall i, k,$$

$$s'^{\text{M}}(D_{ik}) \stackrel{\text{def}}{=} \max\{\alpha_{ik} \cdot s(X) | X \subseteq D_{ik}, X \text{admissible}\}. \quad (12)$$

Obviously we have $s'^{\text{M}}(D_{ik}) = \alpha_{ik} \cdot s^{\text{M}}(D_{ik}) \ \forall i, k$, hence $\alpha_{ik} = 1/s^{\text{M}}(D_{ik})$.

We must now adapt this normalization to the situation where the agents are entitled to different fractions of the resource (non uniform quotas). The way to do this is simple (see for example [Brams and Taylor, 1996, section 2.8]): suppose that we have three agents, with quotas q = (1/10,3/10,6/10). This is equivalent to an equal division between a society of 10 fictitious agents, followed by two groupings of 3 and 6 shares for our last two real agents. This argument leads to an adaptation of the previous normalized individual satisfaction: let state

$$s'^{\text{M}}(D_{ik}) = \alpha_{ik} \cdot s^{\text{M}}(D_{ik}) = q_i \quad (13)$$

instead of 1 as in equation 11, hence $\alpha_{ik} = q_i / s^{\text{M}}(D_{ik})$ To sum up, the set of daily selected pictures with this method maximizes the function

$$s'(A_{\bullet k}) = \sum_{i=1}^{n} q_i \cdot \frac{s(A_{ik})}{s^{\text{M}}(D_{ik})} \quad (14)$$

under admissibility constraints.

With this choice for the coefficients α_{ik}, it is not difficult to see that the selected decisions are independent of the scale of weights used by each agent (in other words, the preference order induced by s' over potential decisions is not changed if some agents multiply their weights by a constant factor). However, the method does not guarantee the satisfaction of the fairness constraint (counter-examples can be easily built). This constraint will have to be checked *a posteriori.* Hopefully, it has a lot of chance to be satisfied, for two reasons :

1. a structural reason : the normalization of the weights tends to favor agents with upper quotas, in a direction favorable to the satisfaction of the fairness constraint; moreover, the fairness constraint is rather soft;
2. a statistical reason : when there is a large number of candidate pictures, not too tightly incompatible, the structural reason can exert its influence; this is the case with our (realistic) simulation data: the simulations show that the fairness constraint is always widely satisfied (see results in section 6).

The function maximized being a linear combination of individual satisfactions of the agents, decisions selected by this method are Pareto-optimal decisions[1] in the n-dirnensional space of individual satisfactions. Such decisions are also called *efficient* decisions. It is impossible to improve a decision selected by this method for one agent without reducing the satisfaction of at least another agent. This property explains the good satisfaction levels obtained with this method in our simulations and justifies the name "efficiency first".

We have designed a variant of this method, for the case where the fairness constraint would not be satisfied, when requests are poorly distributed and highly incompatible. This variant is inspired by the classical Knaster's procedure of sealed bids [Brarns and Taylor, 1996, section 3.2], [Young, 1994, section 8.2]. We compute each day fictitious monetary compensations between agents, reflecting the gap between the actual and ideal shares. An agent having a positive credit is "late" on its quota (it received not enough pictures selected) and conversely, an agent with a negative credit is "ahead" on its quota. These compensations are used to modify the above normalization procedure for the next days in a direction favorable to a fairest share.

This method and its variant can be implemented successfully using the same Valued CSP framework as before. However, the number of instances to be solved is large (all the $s^M(D_{ik})$ must be computed) and the size of the whole instance (for the maximization of $S'(A_{*k})$) may be very important. Our simulations show that an optimal decision can be computed almost all days in a reasonable amount of time. For very large instances, we have to turn to local search procedures (descent search or simulated annealing).

5 Compromises between fairness and efficiency : a multi-criteria approach

The third approach does not focus on fairness or efficiency, but computes a set of good compromise decisions. The aim is to help a human decision-maker to take decisions, by providing this decision-maker with interesting compromises.

The most precise way to set the whole problem is to formulate it as a sequence of multi-criteria discrete optimization problems. The criteria to be maximized would be:

- the *n* agent's satisfaction criteria $cs_i, i = 1, \ldots, n$;
- a criterion *j* measuring the quality of share, to be defined.

Only the set of Pareto-optimal decisions in this *n* + 1 dimensional space are worth considering. The approach which would consist of collecting this set of decisions is unworkable, because it is very large (in our application). A straightforward idea is to select the fairest decision within the set of efficient decisions (see for example [Moulin, 1988, page 14]). It is as well unworkable because the number of potential decisions is too large to allow exhaustive search.

So, we have to resign ourselves to aggregate some criteria. A sensible solution is to aggregate individual satisfactions into a global cumulative satisfaction *ges,* and to keep apart the quality of share criterion *j.* Eventually, potentially interesting decisions will be presented in the two-dimensional space *j* x *yes.*

5.1 Measuring the quality of share

It is questionable to base the quality of share upon the individual satisfactions obtained by agents, because these satisfactions are not expressed in a common scale, and hence are difficult to compare. A better idea is to base our measure upon some function of the *real cost* of pictures, such as time, memory or power consumption on board. Let *c(x)* be the cost of the picture *x.* The cost function is supposed to be independent of the agent requesting the picture, and to have been fixed by mutual agreement between agents. Let

$$c(A_{ik}) \stackrel{\text{def}}{=} \sum_{x \in A_{ik}} c(x), \tag{15}$$

$$cc_i \stackrel{\text{def}}{=} \sum_{k \in I} c(A_{ik}), \tag{16}$$

$$\text{and } \mathbf{cc} \stackrel{\text{def}}{=} (cc_i)_{i=1}^{n}. \tag{17}$$

The last quantity is just the vector of cumulative costs of pictures selected for the agents over the interval /. We propose to measure the quality of share over / by a "distance" between cc and q, the quota vector.

Microeconomists have developed a rich set of *inequality indices* (see for example [Moulin, 1988, section 2.6]), that we can use to base our function *j* measuring the quality of share. The popular Gini indiee

$$G(\mathbf{u}) \stackrel{\text{def}}{=} \frac{1}{2n^2\overline{u}} \sum_{1 \le i,j \le n} |u_i - u_j| \tag{18}$$

measures the inequality resulting from a vector of utilities u = $(u_1, \ldots, u_n)$. *u* is the average value of the U_i. It can be generalized to the non-uniform case to fit our needs, using an argument similar to the one given in section 4, in the following way:

$$G'(\mathbf{u}, \mathbf{q}) \stackrel{\text{def}}{=} \frac{1}{2} \sum_{1 \le i,j \le n} |\tilde{u}_i \cdot q_j - \tilde{u}_j \cdot q_i| \tag{19}$$

[1]A Pareto-optimal decision always heats any other decision on at least one criterion.

$$\text{with } \tilde{u}_i = \frac{u_i}{\sum_{j=1}^{n} u_j}. \tag{20}$$

Taking

$$j \stackrel{\text{def}}{=} 1 - G'(\mathbf{cc}, \mathbf{q}) \tag{21}$$

finishes the job. We have $0 < j < 1$, and $j = 1$ when the share is perfect (costs of obtained pictures exactly proportional to quotas).

5.2 Aggregating individual satisfactions

As a measure of the global cumulative satisfaction of agents over the interval J, we choose a linear combination of normalized cumulative individual satisfactions :

$$gcs(A_{\bullet k}) \stackrel{\text{def}}{=} \frac{1}{n} \sum_{i=1}^{n} \frac{cs_i}{cs_i^{\mathrm{M}}} \tag{22}$$

It has the following properties : $0 < gcs < 1$ (the maximum 1 is reached when each agent is satisfied as much as it can be if it were the only owner of the resource); *gcs* is independent of the individual scales of weights; it is independent of quotas[2].

5.3 Computing decisions

This method is very costly in term of computational resource. The set of Pareto-optimal decisions in the *j* x *gcs* space can be computed exactly by a branch-and-bound search, or approached by an adapted local search method when the search space is too large.

6 Simulations

We have used data from the simulated demand concerning the future Spot5 satellite, which will carry three cameras on board. This data, provided originally for the mono-agent case, has been adapted to simulate a demand from $n = 3$ agents. Simulated agents request each day about the same number of pictures. The quota vector for the simulation is (0.1,0.3,0.6). Weights are in the range 1 to 100. We dispose of data for 371 days. The most loaded day comprises 427 requested pictures and 18878 binary and ternary admissibility constraints. The cost function is simply *c(x)* — 1, Vx (that is, we only count the number of selected pictures). The interval of days / on which cumulative satisfaction and cost functions are based is always / = [1 . . . K], where *k* is the present day.

The table 2 sums up the numerical results obtained from the simulation. It gives the cumulative satisfaction and cost of pictures obtained by agents with each method over the whole simulation interval / = [1...371]. Cumulative satisfactions c_{si} should be compared with the maximal possible cumulative satisfactions *cs* for each agent, given on the second line.

[2]This option is questionable but seems rather sensible, because we consider that the satisfactions of agents are of equal importance, even if they are entitled with different rights. Note that the quota vector is taken into account by the quality of share.

	agent i	1	2	3
	cs_i^{M}	112499	112210	115860
Fairness First	cs_i	13196	34078	66819
	cc_i	585	1750	3405
	cc_i (%)	10.2	30.5	59.3
Efficiency First	cs_i	36491	58036	100408
	$q_i \cdot cs_i^{\mathrm{M}}$	11250	33663	69516
	cc_i	1725	3268	5146
	cc_i (%)	17.0	32.2	50.8
Multi-Criteria ($\lambda = 9$)	cs_i	33928	54075	90472
	cc_i	1296	2773	5156
	cc_i (%)	14.0	30.1	55.9

Table 2: Simulation results for the three methods.

For the Efficiency First method, we give the minimal fair shares $q_i \cdot cs_i^{\mathrm{M}}$. As it can be seen, the fairness constraint (equation 6) is widely satisfied.

For these results, we have simulated a restricted form of the Multi-Criteria method : instead of building a complete set of non-dominated decisions in the *j x gcs* plane, we only look for a decision close to the line of slope A = 9 from the (1,1) point (see figure 1).

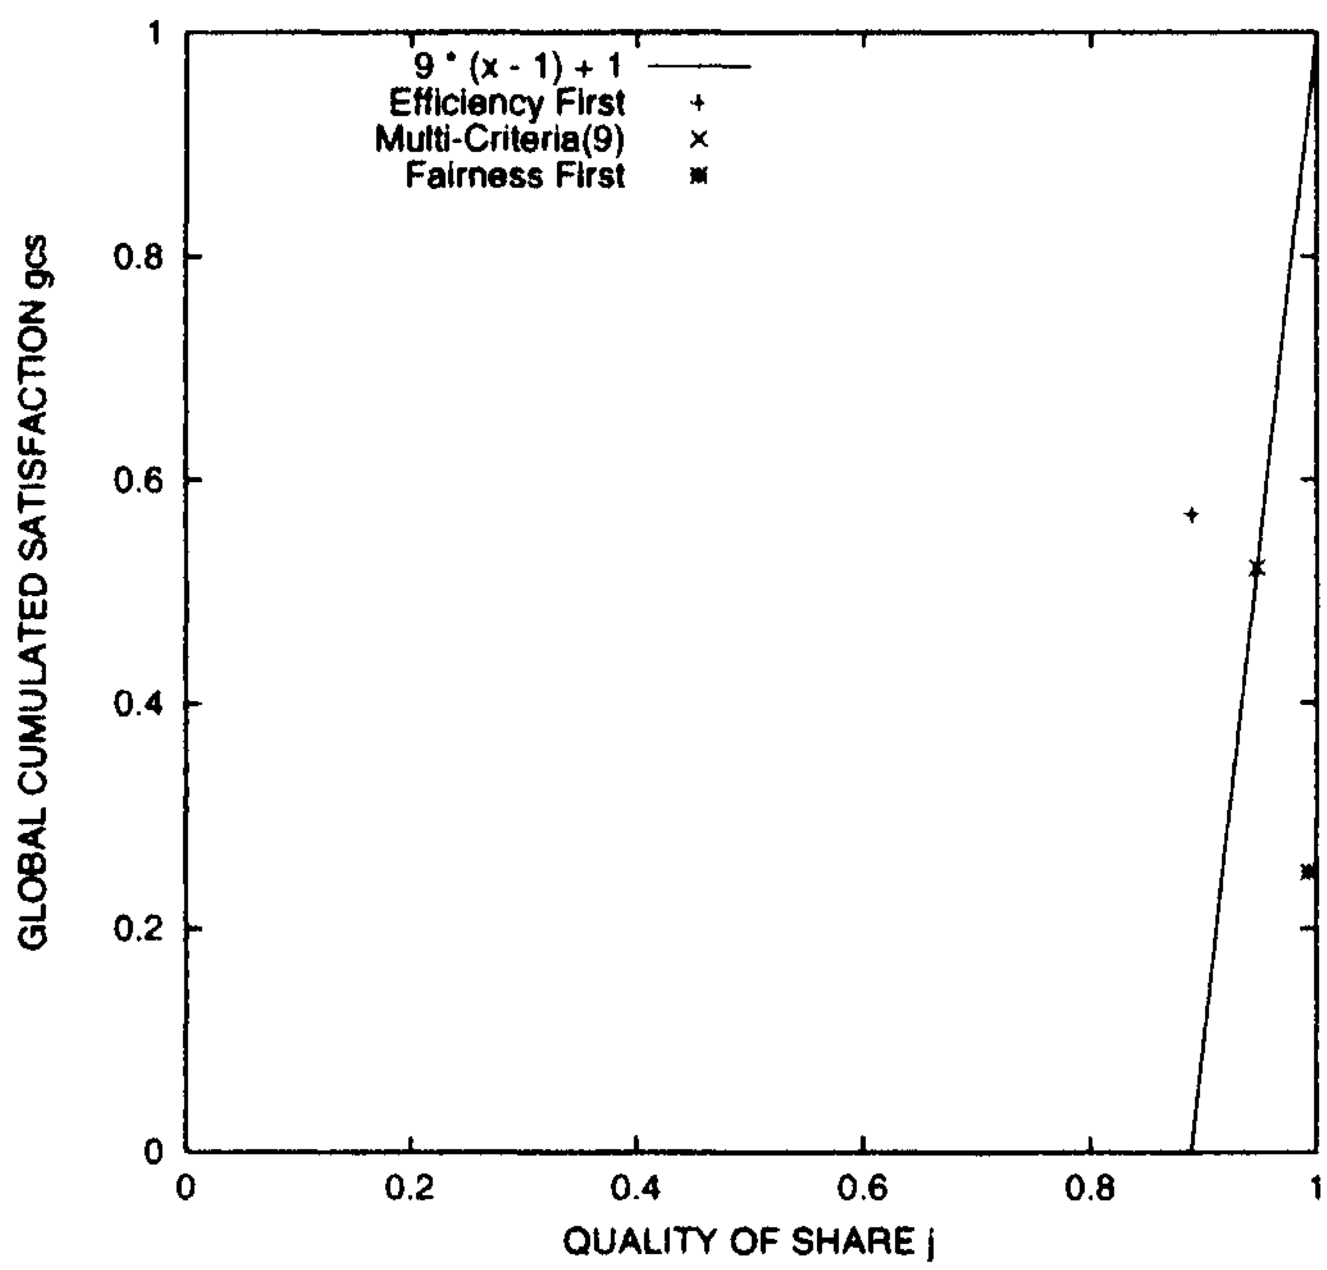

Figure 1: Comparison of methods in the *j* x *gcs* plane.

Finally, the figure 1 sets our three methods in the 2-dimensional plane (quality of share, global cumulative satisfaction). On these two criteria, no method dominates over another. Fairness First provides a quite perfect share, but a poor satisfaction. Efficiency First gives the best satisfaction, but a price in quality of share must be paid for it (this is quite acceptable, since the fairness constraint is satisfied). Lastly, Multi-Criteria($\lambda = 9$) gives a compromise solution between the two others. Other values for λ would allow to get other compromises : this is also of interest in this method.

7 Summary and conclusions

We have described a specific share decision problem involving multiple agents, in which the satisfaction of two kinds of constraints poses a dilemma: efficiency constraints aim at satisfying the agents the most, whereas a fairness constraint watches over equity among agents.

We proposed three different methods to solve this problem. The first method searches for fairness first, and then for efficiency. It is a simple *a priori* sharing method, allocating observation windows to each agent in turn.

The second method is based on the opposite view : first efficiency, fairness if possible. A global satisfaction criterion is defined and maximized. A "minimal fair share" for each agent is defined *a priori* but only checked *a posteriori.*

The third approach does not favor one constraint or the other, but computes a set of good compromise decisions. This is a multi-criteria approach, based on the computation of the set of Pareto-optimal decisions in the two-dimensional space (global-satisfaction, quality-of-share). This set is computed exactly by a branch-and-bound search, or approached by an adapted local search method when the search space is too large.

These three methods have been simulated on the basis of the expected data for the future Spot5 satellite. In short:

- the first method results in very good shares, but inefficient decisions,
- the second one delivers quite good decisions (minimal fair shares are always achieved and the global satisfaction is high), and uses a tolerable amount of computational resources,
- the last one is very costly in computational resources, but allows a human decision-maker to preview a set of interesting non-dominated compromise decisions.

The overall conclusions of this work are:

- no method can be indisputably put forward; the problem is not to choose a method against another one, it is to present to the agents a set of methods and their properties and to let them decide according to the properties they consider the most important[3];
- whereas general methods of sharing can be stated, each share problem is specific and must be studied carefully;
- discrete share problems like this one are computationally very consuming; more specialized combinatorial optimization algorithms are needed to solve them.

[3]See [Rosenschein and Zlotkin, 1994] for a discussion about this point. What we call a method is called by them a

Acknowledgments

The authors would like to thank the anonymous reviewers for their useful comments, as well as Lionel Lobjois and Loi'c Tregan for their help in the simulation work.

References

[Bistarelli *et a/.*, 1995] S. Bistarelli, U. Montanari, and F. Rossi. Constraint solving over semirings. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence,* pages 624630, Montreal, Quebec, Canada, August 1995. International Joint Commitee on Artificial Intelligence.

[Brams and Taylor, 1996] S. J. Brams and A. D. Taylor. *Fair Division — From cake-cutting to dispute resolution.* Cambridge University Press, 1996.

[Freuder and Wallace, 1992] E. Freuder and R. Wallace. Partial Constraint Satisfaction. *Artificial Intelligence,* 58(1-3):21 70, December 1992.

[Keeney and Raiffa, 1976] R. L. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs.* John Wiley and Sons, 1976.

[Larrosa *et al,* 1998] J. Larrosa, P. Messeguer, T Schiex, and G. Verfaillie. Reversible DAC and Other Improvements for Solving Max-CSP. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence,* pages 347-352, Madison, Wisconsin, July 1998. American Association for Artificial Intelligence.

[Moulin, 1988] H. Moulin. *Axioms of Cooperative Decision Making.* Cambridge University Press, 1988.

[Moulin, 1995] H. Moulin. *Cooperative Microeconomics, A Game-Theoretic Introduction.* Prentice Hall, 1995.

[Rosenschein and Zlotkin, 1994] J. S. Rosenschein and G. Zlotkin. *Rules of Encouter.* MIT Press, Cambridge, Massachusetts, 1994.

[Schiex *et al,* 1995] T. Schiex, H. Fargier, and G. Verfaillie. Valued Constraint Satisfaction Problems: Hard and Easy Problems. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence,* pages 631-637, Montreal, Quebec, Canada, August 1995. International Joint Commitee on Artificial Intelligence.

[Verfaillie *et al,* 1996] G. Verfaillie, M. Lemaitre, and T. Schiex. Russian Doll Search for Solving Constraint Optimization Problems. In *Proceedings of the Thirteen National Conference on Artificial Intelligence,* pages 181-187, Portland, Oregon, August 1996. American Association for Artificial Intelligence.

[Wallace, 1994] R. Wallace. Directed Arc Consistency Preprocessing. In *Constraint Processing,* pages 121—137. Lecture Notes in Computer Science 923, Springer, 1994.

[Young, 1994] H. P. Young. *Equity in Theory and Practice.* Princeton University Press, 1994.