ALLOCATION OF DOWNLINK WINDOWS FOR A CONSTELLATION OF SATELLITES

Cédric Pralet[1], Gérard Verfaillie[1], Xavier Olive[2], Serge Rainjonneau[2], and Isabelle Sebbag[3]

[1]*ONERA – The French Aerospace Lab, F-31055, Toulouse, France*
[2]*TAS, Thales Alenia Space, Toulouse, France*
[3]*CNES, Toulouse, France*

ABSTRACT

This paper considers the problem of allocation of downlink windows for a constellation of satellites whose mission is to perform a regular surveillance of ship movements over all the oceans. This surveillance is realized using both radar acquisitions of specific areas and acquisition flows continuously collected over the whole orbits. These acquisitions must be downloaded as quickly as possible, to reduce the so-called *age of information* and improve situation awareness on ground. The paper presents a mathematical modeling of the downlink windows allocation problem in this context, problem decomposition techniques that reduce complexity, and some first experiments on practical instances.

Key words: Downlink, constellation, age of information, mixed integer programming, problem decomposition.

1. PROBLEM PRESENTATION

We consider a problem of allocation of downlink windows for a space mission called SAMSON, whose goal is the regular surveillance of ship movements over all the oceans. This mission is realized using a constellation of four to six satellites. Each satellite has the same capabilities and can perform several kinds of acquisitions:

1. high resolution radar acquisitions of specific areas, called *HR acquisitions*; these acquisitions are mandatory and concern areas of greatest interest;
2. radar acquisitions called *SURV acquisitions*, for regular surveillance of the oceans; these acquisitions are optional and are performed using a radar instrument which can be switched on/off at some times;
3. acquisitions continuously collected over the whole orbits, called *AIS* and *Collect* acquisitions, performed using instruments which are continuously active and capture signals emitted by ships.

HR acquisitions correspond to individual data, while non-HR ones (SURV, AIS, Collect) correspond to data flows. Acquisitions recorded on-board can be downloaded using a certain number of *visibility windows* of ground stations. A satellite cannot simultaneously download data to two ground stations, and one ground station cannot receive data from two satellites in parallel. Moreover, at the level of each station, there is a certain duration necessary for reconfiguration between the tracking of two successive satellites. The problem is then to build a conflict-free allocation of *actual downlink windows* to each satellite. These downlink windows must be placed in existing visibility windows of ground stations. The objective is to allocate downlink windows once and for all in preprocessing by estimating possible downloads, and not to decide on precise downloads to be performed.

Another aspect is that the allocation produced must be such that data is downloadable as quickly as possible, in order to improve situation awareness on ground. In this direction, a key criterion is the age of each acquisition, defined as the temporal distance between the realization of an acquisition and its download. Minimizing the age of HR acquisitions is also considered as a prioritary objective with regard to the other kinds of acquisitions.

The paper is organized as follows. Section 2 defines the problem data, Section 3 gives a mathematical modeling of the problem, Section 4 introduces problem decomposition techniques that reduce complexity, Section 5 details resolution techniques, Section 6 provides some experimental results, and Section 7 discusses related works.

2. PROBLEM DATA

We consider a set of satellites S, a set of reception ground stations R, and a set of visibility windows W. Each visibility window $w \in W$ is defined as a tuple $(S(w), R(w), Ts(w), Te(w))$, with $S(w)$ and $R(w)$ the satellite in S and the ground station in R associated with w, and $Ts(w)$, $Te(w)$ the start and end times of w.

Given a satellite s, $W(s)$ denotes the sequence of visibility windows $w \in W$ such that $S(w) = s$, ordered by increasing start times. $|W(s)|$ denotes the length of this sequence, and $W(s)$ is written as $[W(s,1), \ldots, W(s, |W(s)|)]$. For each $i \in [1..|W(s)|]$,

$Ts(W(s,i))$ and $Te(W(s,i))$ are denoted more concisely as $Ts(s,i)$ and $Te(s,i)$ respectively. Two distinct visibility windows $w, w' \in W(s)$ such that w appears before w' in $W(s)$ are said to be in conflict iff they overlap ($Ts(w') < Te(w)$). On the example of Fig. 1, $W(s_1) = [w_1, w_2, w_3]$ and there is a conflict between windows w_1 and w_2. The end time of the visibility window associated with s that immediately precedes the first visibility window in $W(s)$ is denoted $Te^-(s)$. The memory size occupied by non-HR acquisitions (SURV, AIS, Collect) at that time is denoted $Me^-(s)$. The start time of the visibility window that immediately follows the last visibility window in $W(s)$ is denoted $Ts^+(s)$.

Given a ground station $r \in R$, $W(r)$ denotes the sequence of visibility windows $w \in W$ such that $R(w) = r$, ordered by increasing start times. The duration required by r for station reconfiguration between two successive satellite tracking is denoted $\Delta(r)$ (with $\Delta(r) > 0$). Two distinct visibility windows $w, w' \in W(r)$ such that w appears before w' in $W(r)$ are said to be in conflict iff the temporal distance between w and w' is less than reconfiguration time $\Delta(r)$ ($Ts(w') < Te(w) + \Delta(r)$). On the example of Fig. 1, $W(A) = [w_1, w_5]$ and visibility windows w_1 and w_5 are in conflict for the access to station A. The set of pairs of windows (w, w') in $W(r)$ that are in conflict with each other is denoted $Wc(r)$.

A visibility window $w \in W$ is said to be *conflict-free* iff it has no conflict with other windows, both in terms of satellites and stations. The set of conflict-free visibility windows is denoted Wnc. In Fig. 1, $Wnc = \{w_3, w_4\}$.

The set of mandatory HR acquisitions to be performed is denoted H. The end time associated with the realization of h and the duration required for downloading h are denoted $Te(h)$ and $DuDl(h)$ respectively. The satellite that must realize $h \in H$ is known beforehand and is denoted $S(h)$. Given a satellite s, $H(s)$ denotes the set of HR acquisitions $h \in H$ to be realized by s ($S(h) = s$).

In the mission specification, there also exists a minimum duration $Dmin$ for allocated downlink windows (with $Dmin > 0$). The download rate common to every satellite in S is denoted $DlRate$ ($DlRate > 0$). The mean acquisition rate associated with non-HR acquisitions, also common to all satellites, is denoted $AcqRate$ (with $AcqRate > 0$). $AcqRate$ can be seen as the estimated mean rate at which data flows generated by SURV, AIS, and Collect acquisitions are received on-board.

3. PROBLEM MODELING

3.1. Simplified model

To simplify the presentation, we first consider that each visibility window must contain exactly one allocated downlink window. This assumption will be relaxed in the full model given later in Section 3.2. We also assume that

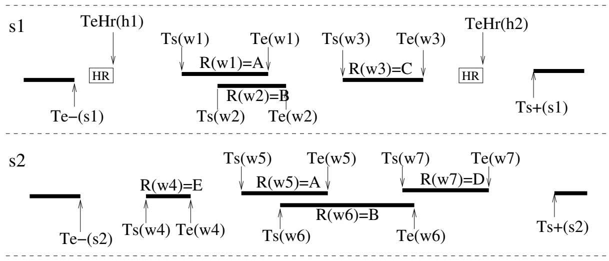

Figure 1. Illustration of problem data

each satellite s uses visibility windows in $W(s)$ chronologically. This means that given two visibility windows $w_1, w_2 \in W(s)$ such that w_1 precedes w_2 in $W(s)$, the downlink window allocated in w_1 must precede the downlink window allocated in w_2. This earliest download heuristics at the level of each satellite simplifies the modeling. On the contrary, there is no restriction on the order between downlink windows that must share a given station, since a station may prefer waiting for a satellite having more prioritary acquisitions on-board.

Decision variables For each visibility window $w \in W$, we define two continuous decision variables $ts(w), te(w) \in [Ts(w), Te(w)]$, that respectively represent the start and end times of the downlink window allocated in w. The duration of this downlink window is $du(w) = te(w) - ts(w)$. For model readability reasons, given a window $W(s,i)$ in the sequence of visibility windows of satellite s, $ts(W(s,i))$ and $te(W(s,i))$ are also denoted as $ts(s,i)$ and $te(s,i)$. In addition, we consider that $te(s,0)$ and $ts(s,|W(s)|+1)$ are defined and take values $Ts^+(s)$ and $Te^-(s)$ respectively.

We also introduce, for each reception station $r \in R$ and for each pair of conflicting visibility windows $(w, w') \in Wc(r)$, one variable $b(w,w') \in \{0,1\}$ specifying whether the downlink window allocated in w precedes the one allocated in w' (value 1) or not (value 0).

Constraints Constraint 1 imposes a minimum duration for downlink windows. Constraint 2 ensures a maximum use of conflict-free visibility windows. Constraint 3 guarantees no overlapping between successive downlink windows associated with the same satellite. Constraints 4 and 5 impose a sufficient transition time between conflicting downlink windows associated with the same station.

$$\forall w \in W \; : du(w) \geq Dmin \tag{1}$$

$$\forall w \in Wnc : (ts(w) = Ts(w)) \wedge (te(w) = Te(w)) \tag{2}$$

$$\forall s \in S, \forall i \in [2..|W(s)|], \; Ts(s,i) < Te(s,i-1) : \quad ts(s,i) \geq te(s,i-1) \tag{3}$$

$$\forall w, w' \in W \text{ s.t. } \exists r \in R, \; (w,w') \in Wc(r) \; :$$

$$(b(w,w') = 1) \rightarrow (ts(w') \geq te(w) + \Delta(r)) \tag{4}$$

$$(b(w,w') = 0) \rightarrow (ts(w) \geq te(w') + \Delta(r)) \tag{5}$$

Objective u_{hr} The first and prioritary objective function considered, denoted u_{hr}, corresponds to the mean age of HR acquisitions. The *age* of an acquisition a refers to the temporal distance between the end of the realization of a and the start of its download. In order to obtain a definition independent from the precise scheduling of downloads, we build the following approximation. Assume that a ends at time t and is downloaded in a downlink window starting at time t'. Then, (1) if $t < t'$, $age(a) = t' - t$, (2) otherwise $age(a) = 0$. This is equivalent to be optimistic and assume that a is downloaded as soon as possible in its associated downlink window. The two previous cases can be gathered using $age(a) = \max(0, t' - t)$. For acquisitions a not downloaded over the planning horizon and to be performed by satellite s, we take $age(a) = Ts^{+}(s) - t$, that is we do as if a is downloaded as soon as possible in the visibility window that follows the last window in $W(s)$.

In order to define objective function u_{hr}, the following decision variables are added to the model:

- $\forall s \in S$, $\forall i \in [1..|W(s)| + 1]$, $\forall h \in H(s)$, $dlHr(s,i,h) \in \{0,1\}$: indicates whether HR acquisition h is downloaded in $W(s,i)$ (value 1) or not (value 0); for $i = |W(s)| + 1$, $dlHr(s,i,h) = 1$ means that h is downloaded after the end of the planning horizon;
- $\forall h \in H$, $age(h) \in \mathbb{R}^{+}$: age of HR acquisition h;
- $\forall s \in S$, $\forall i \in [1..|W(s)|]$, $duDlHr(s,i) = \sum_{h \in H(s)} DuDl(h) \cdot dlHr(s,i,h)$: duration consumed by the download of HR acquisitions in visibility window $W(s,i)$.

Constraints 6 to 9 are imposed over these variables. Constraint 6 ensures that each HR acquisition is downloaded exactly once (possibly after the end of the planning horizon). Constraint 7 expresses that the download duration in a given downlink window must not exceed the duration of that window. Constraint 8 defines the age of an HR acquisition. Constraint 9 gives an additional temporal condition on the end time of the downlink window in which an HR acquisition is downloaded.

$$\forall s \in S, \forall h \in H(s) : \sum_{1 \leq i \leq |W(s)|+1} dlHr(s,i,h) = 1 \quad (6)$$

$$\forall s \in S, \forall i \in [1..|W(s)|] : duDlHr(s,i) \leq du(s,i) \quad (7)$$

$$\forall s \in S, \forall i \in [1..|W(s)| + 1], \forall h \in H(s) : (dlHr(s,i,h) = 1) \rightarrow (age(h) = \max(0, ts(s,i) - Te(h))) \quad (8)$$

$$\forall s \in S, \forall i \in [1..|W(s)|], \forall h \in H(s) : (dlHr(s,i,h) = 1) \rightarrow (te(s,i) \geq Te(h) + DuDl(h) \quad (9)$$

In theory, there exist particular cases in which these constraints do not suffice to guarantee that the downlink windows chosen for HR acquisitions can actually be used. Nevertheless, the model proposed provides a quite reasonable approximation of the age of HR acquisitions.

The mean age objective function u_{hr} to be minimized is then given by Equation 10. Other criteria such as the worst age could also be considered.

$$u_{hr} = \frac{1}{|H|} \cdot \sum_{h \in H} age(h) \quad (10)$$

Objective u_{nhr} The secondary objective function, denoted u_{nhr}, corresponds to the mean age of non-HR acquisitions. In order to define the age of these acquisitions, the following variables are added to the model:

- $\forall s \in S$, $\forall i \in [0..|W(s)|]$, $me(s,i) \in \mathbb{R}^{+}$: estimation of the memory size occupied by non-HR data at the end of the downlink window placed in $W(s,i)$;
- $\forall s \in S$, $\forall i \in [1..|W(s)| + 1]$, $to(s,i) \in \mathbb{R}$: estimation of the acquisition time of the oldest non-HR data to be downloaded in $W(s,i)$.

These variables must satisfy constraints 11 to 13. Constraints 11 and 12 define the memory size occupied by non-HR acquisitions at the end of each downlink window. In Constraint 12, this memory size takes into account (1) the memory size occupied at the end of the previous downlink window, (2) the data volume generated by data flows (SURV, AIS, and Collect) between the end of the previous downlink window and the end of the current downlink window, and (3) the maximum volume of non-HR data downloadable in the current downlink window. Constraint 13 defines the date of the oldest non-HR acquisition to be downloaded in a given window.

$$\forall s \in S : me(s,0) = Me^{-}(s) \quad (11)$$

$$\forall s \in S, \forall i \in [1..|W(s)|] : me(s,i) = \max(0, me(s,i-1) + AcqRate \cdot (te(s,i) - te(s,i-1)) - DlRate \cdot (du(s,i) - duDlHr(s,i))) \quad (12)$$

$$\forall s \in S, \forall i \in [1..|W(s)| + 1] : to(s,i) = te(s,i-1) - \frac{me(s,i-1)}{AcqRate} \quad (13)$$

The notions are illustrated in Fig. 2. In Fig. 2(a), downlink window number $i - 1$ ends at time $t = 10$ and the memory size occupied at the end of this window is null. This means that all acquisitions performed strictly before $t = 10$ are downloaded before the end of window $i - 1$. The oldest non-HR acquisition to be downloaded in window i therefore corresponds to the acquisition realized just after the end of window $i - 1$, that is $to(s,i) = 10$. In Fig. 2(b), downlink window number $i - 1$ ends at time $t = 9$ and the memory size occupied at that time is 6 memory units. Given that data flows are recorded with rate $AcqRate = 3$ and if we consider that data flows are downloaded chronologically (in the order of their acquisition), then the oldest acquisition in memory at time $te(s, i-1)$ was realized $\frac{me(s,i-1)}{AcqRate}$ units of time before $te(s, i-1)$, hence $to(s, i) = te(s, i-1) - \frac{me(s,i-1)}{AcqRate} = 7$. Fig. 2(c) provides an example showing that $to(s, i)$ can even be located before the start of window $i-1$.

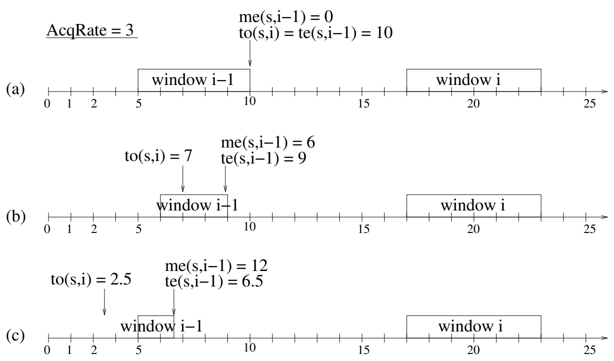

Figure 2. Computation of $to(s, i)$

Similarly to HR acquisitions, the goal is to minimize the mean age of non-HR data. The main difference with HR acquisitions is that we have to deal with data flows instead of data punctually produced by acquisitions of specific areas. To define the mean age of a flow, let us consider a data flow recorded from t_1 to t_2 and downloaded in a window starting at time t'. By reusing the same definition for the age of a single acquisition as previously, the mean age of data recorded from t_1 to t_2 is given by $\frac{1}{t_2-t_1} \cdot \int_{t_1}^{t_2} \max(0, t'-t)\, dt$. By reasoning over the whole horizon, of length $L(s) = Te^-(s) - Ts^+(s)$, the mean age of data flows is given by Equation 14. In this equation, we compute the sum of the ages of data downloaded in each downlink window (terms $\int_{to(s,i)}^{to(s,i+1)} \max(0, ts(s,i)-t)\, dt$). We also consider the particular case of the end of the horizon, by computing the mean age of acquisitions downloaded at $Ts^+(s)$ (term $\int_{to(s,|W(s)|+1)}^{Ts^+(s)} \max(0, Ts^+(s)-t)\, dt$). Via some computations, Equation 14 can be reformulated as in Equation 15.

$$\frac{1}{L(s)} \cdot \left(\sum_{i \in [1..|W(s)|]} \int_{to(s,i)}^{to(s,i+1)} \max(0, ts(s,i)-t)\, dt + \int_{to(s,|W(s)|+1)}^{Ts^+(s)} \max(0, Ts^+(s)-t)\, dt \right) \quad (14)$$

$$= \frac{1}{2L(s)} \cdot \left(\sum_{i \in [1..|W(s)|+1]} (ts(s,i) - to(s,i))^2 - \sum_{\substack{i \in [1..|W(s)|] \\ ts(s,i) > to(s,i+1)}} (ts(s,i) - to(s,i+1))^2 \right) \quad (15)$$

We decide in the following to ignore the second sum in Equation 15, for several reasons: (1) in practice, it is expected that acquisitions recorded on-board are likely to be downloadable in the next downlink window; in this case, $ts(s, i) > to(s, i+1)$ will often be false and the second sum will only contain a few terms; (2) because of the minimum duration of downlink windows and because of the distance between successive visibility windows, it is expected that for most of indices $i \in [1..|W(s)|]$, terms $(ts(s,i) - to(s,i+1))^2$ of the second sum are much smaller than terms $(ts(s,i+1) - to(s,i+1))^2$ of the first sum; (3) discarding the second sum allows to get a convex (quadratic) criterion, which makes the optimization task easier.

As a result, we keep only the first term in Equation 15 and define objective function u_{nhr} to be minimized as:

$$u_{nhr} = \frac{1}{|S|} \sum_{s \in S} \frac{1}{2L(s)} \sum_{i \in [1..|W(s)|+1]} (ts(s,i) - to(s,i))^2 \quad (16)$$

Equation 16 shows that in order to decrease the mean age of acquisitions, it is interesting to minimize the gap between $to(s, i)$ and $ts(s, i)$, that is to increase $to(s, i)$ and decrease $ts(s, i)$. It can be shown that, as a side effect, it indirectly maximizes the duration of downlink windows and minimizes the distance between these windows.

3.2. Complete Model

In the previous model, each visibility window $w \in W$ must contain exactly one allocated downlink window. But in practice, having one downlink window of duration $\geq Dmin$ in every $w \in W$ may be impossible. It may also be suboptimal, due to the time wasted by station reconfiguration between the tracking of two satellites.

This is why we propose a second model, in which each visibility window can contain 0 or 1 downlink window. To allow for 0 to K downlink windows inside a single visibility window (not just 0 or 1), it would suffice to duplicate visibility windows K times, the duplicated windows being in conflict with each other. This can be useful for situations in which one ground station tracks a first satellite, then tracks a second one for receiving an important acquisition, and then comes back to the first one.

Additional variables The main difference in terms of modeling is the addition of variables representing the sequence of allocated downlink windows for each satellite. More formally, the variables added are as follows:

- $\forall w \in W$, $use(w) \in \{0, 1\}$: existence of a downlink window in visibility window w (value 1) or not (value 0); for $s \in S$ and $i \in [1..|W(s)|]$, $use(W(s, i))$ is denoted more concisely $use(s, i)$;

- $\forall s \in S, \forall i < i' \in [0..|W(s)|+1]$, $next(s,i,i') \in \{0,1\}$: indicates whether for satellite s, downlink window placed in $W(s,i)$ is the downlink window that immediately precedes downlink window placed in $W(s,i')$ (value 1) or not (value 0).

Constraints To avoid any confusion, we give in Equations 17 to 27 the set of new constraints considered in this second model. Constraints 17-18 express that a downlink window is used iff it is involved in the sequence of allocated downlink windows (it has exactly one successor and one predecessor). Constraint 19 ensures that the window that ends at $Te^-(s)$ has exactly one successor. Constraint 20 ensures that the window that starts at $Ts^+(s)$ has exactly one predecessor. Constraint 21 gives arbitrary values to variables of precedence between unused windows. Last, Constraints 22 to 27 correspond to a reformulation of Constraints 1 to 5 of the simplified model.

$$\forall s \in S, \forall i \in [1..|W(s)|]:$$
$$use(s,i) = \sum_{i' \in [i+1..|W(s)|+1]} next(s,i,i') \quad (17)$$
$$use(s,i) = \sum_{i' \in [0..i-1]} next(s,i',i) \quad (18)$$
$$\forall s \in S, :$$
$$\sum_{i \in [1..|W(s)|+1]} next(s,0,i) = 1 \quad (19)$$
$$\sum_{i \in [0..|W(s)|]} next(s,i,|W(s)|+1) = 1 \quad (20)$$
$$\forall w, w' \text{ s.t. } \exists r \in R, (w,w') \in Wc(r):$$
$$(b(w,w') \leq use(w)) \wedge (b(w,w') \leq use(w')) \quad (21)$$
$$\forall w \in W : du(w) \geq Dmin \cdot use(w) \quad (22)$$
$$\forall w \in Wnc : (use(w) = 1) \wedge$$
$$(ts(w) = Ts(w)) \wedge (te(w) = Te(w)) \quad (23)$$
$$\forall w \in W \setminus Wnc :$$
$$(use(w) = 0) \rightarrow (ts(w) = te(w) = Ts(w)) \quad (24)$$
$$\forall s \in S, \forall i_1 < i_2 \in [1..|W(s)|], Ts(s,i_2) < Te(s,i_1):$$
$$(next(s,i_1,i_2) = 1) \rightarrow (ts(s,i_2) \geq te(s,i_1)) \quad (25)$$
$$\forall w, w' \in W \text{ s.t. } \exists r \in R, (w,w') \in Wc(r):$$
$$(b(w,w') = 1) \rightarrow (ts(w') \geq te(w) + \Delta(r)) \quad (26)$$
$$((b(w,w') = 0) \wedge (use(w) = 1) \wedge (use(w') = 1))$$
$$\rightarrow (ts(w) \geq te(w') + \Delta(r)) \quad (27)$$

Objective Objectives u_{hr} and u_{nhr} must be reformulated as well. For u_{hr}, the only difference is the addition of Constraint 28 below, which expresses that an HR acquisition cannot be downloaded in a downlink window that is not used. Constraints 6 to 9 are kept.

$$\forall s \in S, \forall i \in [1..|W(s)|], \forall h \in H(s):$$
$$dlHr(s,i,h) \leq use(s,i) \quad (28)$$

The expression of u_{hr} is still that of Equation 10.

As for u_{nhr}, Constraints 11 to 13 of the simplified model are replaced by Constraints 29 to 33. Constraints 29 to 31 redefine the memory size $me(s,i)$ occupied at the end of each downlink window, and Constraints 32-33 redefine the values of variables $to(s,i)$ which give an estimation of the realization time of the oldest non-HR data to be downloaded in $W(s,i)$.

$$\forall s \in S: me(s,0) = Me^-(s) \quad (29)$$
$$\forall s \in S, \forall i \in [1..|W(s)|], W(s,i) \notin Wnc:$$
$$(use(s,i) = 0) \rightarrow (me(s,i) = 0) \quad (30)$$
$$\forall s \in S, \forall i < i' \in [0..|W(s)|]: (next(s,i,i') = 1) \rightarrow$$
$$(me(s,i') = \max(0, me(s,i)$$
$$+ AcqRate \cdot (te(s,i') - te(s,i))$$
$$- DlRate \cdot (du(s,i') - duDlHr(s,i')))) \quad (31)$$
$$\forall s \in S, \forall i \in [1..|W(s)|], W(s,i) \notin Wnc:$$
$$(use(s,i) = 0) \rightarrow (to(s,i) = Ts(s,i)) \quad (32)$$
$$\forall s \in S, \forall i < i' \in [0..|W(s)|+1]:$$
$$(next(s,i,i') = 1) \rightarrow$$
$$(to(s,i') = te(s,i) - \frac{me(s,i)}{AcqRate}) \quad (33)$$

The expression of u_{nhr} is still the expression given in Equation 16 (in particular thanks to the fact that the constraints ensure that $ts(s,i) = to(s,i) = Ts(s,i)$ for a window $W(s,i)$ which is not used).

4. PROBLEM DECOMPOSITION

The model previously described contains a certain number of variables and constraints. Using this model over a large horizon may lead to high computation times. This is why we propose a problem decomposition technique that identifies small sets of conflicting windows and allows to reason locally on each of these sets. These elementary sets of conflicting windows are obtained via the notions of conflict graph and conflict group introduced below.

Conflict graph and conflict groups The *conflict graph* is the graph which contains one node per visibility window $w \in W$, and whose edges are defined inductively as follows:

1. if there is a conflict between visibility windows w and w' (in terms of satellite or station), then there is an edge between w and w';
2. given a satellite s, if windows $w = W(s,i)$ and $w' = W(s,i+1)$ have conflicts with other windows (but not necessarily with each other), then there is an edge between w and w';
3. if a visibility window $W(s,i)$ is conflict-free but if $W(s,i-1)$ and $W(s,i+1)$ are linked by a path

in the conflict graph, then there exists one edge between $W(s, i-1)$ and $W(s, i)$, and one edge between $W(s, i)$ and $W(s, i+1)$;

4. all edges are obtained using the three previous rules.

Informally, a conflict graph indicates whether decisions concerning two visibility windows w, w' depend on each other, either due to constraints (case 1 above), or due to the criterion (case 2 and 3 above). An example of conflict graph is given in Fig. 3. In this graph, edges are labeled with the rule that created them.

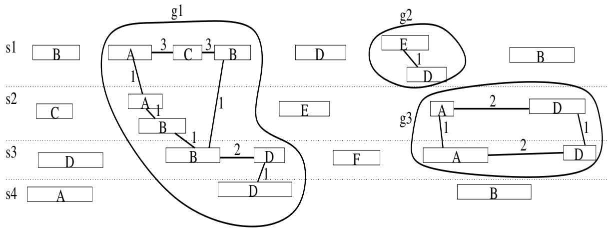

Figure 3. Conflict graph and conflict groups on an example involving 4 satellites (s_1 to s_4) and 6 ground stations (A to F); conflict groups not reduced to a single visibility window are encircled

We define the *set of conflict groups* as the set of connected components of the conflict graph, that is as the set of maximum subgraphs of this graph such that two nodes in the subgraph are connected to each other by a path. The size of a conflict group g is the number of visibility windows it contains. The conflict graph of Fig. 3 contains three connected components (three conflict groups) not reduced to a single node. Extracting conflict groups by computation of connected components has a linear time complexity in the number of nodes and edges in the conflict graph.

The problem decomposition technique proposed consists in reasoning conflict group by conflict group, in order to deal with problems of smaller size. The kind of constellation considered for the Ocean surveillance mission (4 to 6 satellites on low-altitude orbits) implies that group of conflicts are small in practice.

The theoretical justification of the decomposition proposed is that if we assume that it is always possible to download the whole memory during a conflict-free visibility window, then no constraint links variables of a conflict group to variables of another one, and the objective function can be split into several independent local objective functions (one per conflict group).

In case dumping the whole memory during one conflict-free visibility window is not possible, the approach can still be used by solving conflict groups chronologically. This allows to determine the set of HR acquisitions not downloaded yet for a given conflict group, and to determine the memory size occupied by non-HR data at the beginning of each conflict group (element $Me^-(s)$). Note that such a chronological conflict group resolution strategy offers no optimality guarantee on the solution found.

5. RESOLUTION TECHNIQUES

Once the problem is decomposed, several approaches can be used to reason on each conflict group independently:

- Using the model of Section 3.2 and exact resolution techniques on the elementary problems obtained may be possible. All constraints of the model are linear or can be linearized. Objective function u_{hr} is linear, while objective function u_{nhr} is quadratic. Variables are either boolean or continuous. The mathematical problem obtained is therefore a *Mixed Integer Linear Programming* (MIP) problem when considering objective function u_{hr} and a *Mixed Quadratic Programming* (MQP) problem when considering objective function u_{nhr}. Several tools can be used to solve MIP problems, such as IBM ILOG CPLEX. The latter is also able to solve MQP problems under some restrictions that are satisfied by objective function u_{nhr}. To deal with the two objective functions in a hierarchical way, it suffice to solve the MIP problem involving u_{hr}, and then to solve the MQP problem involving u_{nhr}, with the additional constraint that u_{hr} must be equal to the optimal value found at the first step.
- Local search techniques can also be used. In this context, local moves can correspond to additions and removals of downlink windows. At any step, the quality of the current allocation of downlink windows can be evaluated by solving a pure continuous linear/quadratic program, in which the sequence of downlink windows to be used is known (but not their start and end times). In a greedy version, only additions of downlink windows are used.
- A decision rule DR applied chronologically can be defined to solve conflicts. (a) In case of overlapping between two visibility windows associated with the same satellite, we first try to entirely select the window that has the earliest start time and to prune the second window accordingly; if the length of this second pruned window is less than $Dmin$, the opposite choice is made (full selection of the second window and pruning of the end time of the first window); if both options are impossible, the longest window is chosen. (b) In case of a conflict of access to the same station r between visibility windows w_1, w_2 associated with two satellites, let d be the length of the time interval that covers w_1, w_2 ($d = \max(Te(w_1), Te(w_2)) - \min(Ts(w_1), Ts(w_2))$).
  - If $d \leq \Delta(r)$, duration d is too small for using the two windows, hence only one visibility window is kept. In version DR1, the window is selected for the satellite with the oldest preceding downlink window. In version DR2, the window selected is chosen so as to minimize the maximum duration during which no downlink is possible (DR2 considers both the previous and the next possible downlink windows).

- Otherwise ($d > \Delta(r)$), if we assume that w_1 has the earliest start time, then the first allocated window is put in w_1. It starts at $Ts(w_1)$ and ends at $te_1 = \min(Ts(w_1)+(d-\Delta(r))/2, Te(w_1))$ (fair sharing of the available communication time). The second allocated window, put in w_2, starts at $te_1 + \Delta(r)$ and ends at $Te(w_2)$.

6. EXPERIMENTS

A few experiments were performed using IBM ILOG CPLEX. We considered the configuration of the SAMSON constellation involving only four satellites. Over a temporal horizon of 28 days (cycle of the constellation), there are 98 conflict groups corresponding to conflict of access to a ground station. All these conflict groups are of size 2, which can be explained by the kind of constellation considered: two orbital planes; circular, low-altitude, quasi-polar orbits; two satellites per orbital plane, at symmetric positions. A conflict can actually exist only if one satellite of the first orbital plane and one satellite of the second one are visible at the same time from the same ground station. Such conflicts appear only for stations near to the pole. The size of conflict groups could be higher for more complex constellations or for constellations with medium- or high-altitude orbits, for which potential communication windows are longer.

Table 1 gives, for three of the conflict groups obtained, the mean age of acquisitions given by objective u_{nhr} (no HR acquisition here). Two configurations in terms of min duration $Dmin$ of a downlink window and in terms of reconfiguration time $\Delta(r)$ for a station are considered. The results show that on conflicts of size 2, decision rules DR1/DR2 are quite good compared to the optimal MQP strategy, with a small advantage to DR2. Higher gaps between MQP and the decision rules could however be observed on more complex configurations. Concerning computation times, the optimal MQP approach takes about 0.01s in each case. It would be interesting to see on more complex practical instances the evolution of computation times with the size of conflict groups. The techniques of allocation of downlink windows were integrated in a more general tool for managing the SAMSON mission (see [17]), based on decision rule DR1.

		u_{nhr} (in sec.)		
conflict	$Dmin$ / $\Delta(r)$	MQP	DR1	DR2
c1	300s / 600s	**1478.7**	**1478.7**	**1478.7**
c2	300s / 600s	**1087.9**	1096.0	**1087.9**
c3	300s / 600s	**1498.9**	**1498.9**	**1498.9**
c1	100s / 120s	**1325.5**	1334.6	1334.6
c2	100s / 120s	**944.3**	954.9	954.9
c3	100s / 120s	**1334.7**	1340.6	1340.6

Table 1. Comparison of different approaches for some conflict groups

7. RELATED WORKS

One basic issue when dealing with spacecraft sharing ground stations and ground antennas is the allocation of these resources to spacecraft over time. As stated in [15], such an issue has already been tackled using automatic tools for several space missions. Several existing contributions are listed below.

Air Force Satellite Control Network As far as we know, the automatic allocation of communication windows to several satellites was first considered for the Air Force Satellite Control Network (AFSCN). As described in [1], this network contains 16 antennas over 9 ground stations, and receives each day approximately 500 requests of communication windows (requests of communication of a given duration with some antennas, in a time interval defined by an earliest communication time, a latest communication time). The goal is to produce an allocation of communication windows so as to minimize the number of unsatisfied requests. Several approaches were proposed to solve this problem, from the first approaches using MIP and heuristic search [6, 14], to later more efficient approaches based on genetic algorithms [13, 1].

Deep Space Network Another important communication network that led to many works on the topic is the Deep Space Network (DSN). This network is essentially used for deep space missions and missions beyond geostationary orbits. As it is described in [10], it is made of 13 antennas over 3 geographical sites. In addition to antennas, each site also contains equipments that must be shared between antennas. The network is used by 35 users. Similarly to the AFSCN, several approaches were proposed for solving the ground equipment allocation problem, from techniques in [2] based on a MIP formulation, to local search techniques [3, 4], systematic tree search, and genetic algorithms [7]. Recent approaches allow high-level communication requests to be defined [11], interaction with the network users [10], and assistance to conflict resolution [9].

ESA Tracking network Another example is the ESA Tracking network (ESTRACK). As described in [5, 8], this network is composed of 9 stations hold by ESA and of 3 external stations. These stations are used by 10 ESA missions and by some external users. One specificity is the great variability of the missions considered in terms of duration of possible communication windows and in terms of number of communications required over a given horizon. The requirements may correspond to a given periodicity of access to a ground station, minimum and maximum durations of a communication window, minimum and maximum distances between two such windows, or even more complex constraints. The first approach defined in [5] for managing the ESTRACK consists in building a valid allocation of downlink windows for the different missions (no criterion). To do this, an incremental approach that heuristically selects and adds at each step a non satisfied communication request is used, as well temporal constraint reasoning techniques. Optimization of the ESTRACK was added in [8], where techniques based on dynamic programming and local search were explored.

Academic ground stations network [15] considers another kind of network, allowing academics to handle small satellites (CubeSat). In this context, a priority is associated with each communication request and an important aspect concerns the redundancy and fair sharing of communications among the different users [16]. Another specificity is that the number of satellites on the same orbit can be high. The two allocation techniques proposed in [15] are based on branch and bound and gradient descent respectively.

TerraSAR-X/TanDEM-X The allocation of communication windows is also implicitly present in specific missions such as TerraSAR-X and TanDEM-X [12]. The latter are realized thanks to two satellites, TSX-1 and TDX-1, evolving at a distance of some hundreds of meters. Depending on the actual distance between the two satellites, data download in parallel may be impossible, due to interferences. Satellites can however never simultaneously dump data to the same antenna. If TSX-1 and TDX-1 are near, then the time required for station handover is reduced. Techniques described in [12] for managing data downloads is a kind of well-informed greedy search, that chooses at each step to download a given acquisition over the best group of stations. The heterogeneity of satellites TDX-1/TSX-1 in terms of memory available on-board is also handled to obtain a fair sharing of communication windows.

Comparison with the SAMSON mission Compared to all these missions, the formal model proposed in this paper can be related to MIP-based approaches of the literature cited above. However, the SAMSON mission considered in this paper presents several specificities that prevented us from directly reusing existing approaches. One of the main specificities concerns the optimization of (1) the age of acquisitions of specific areas and (2) the age of data generated by acquisition flows. This notion of age of acquisition and acquisition flows is crucial for the ocean global surveillance mission, and may be reused for future missions. Another specificity is the shape of the SAMSON constellation, which makes the problem decomposition approach very efficient.

REFERENCES

[1] Barbulescu, L., Watson, J.-P., Whitley, L. D., and Howe, A. E. (2004). Scheduling space-ground communications for the air force satellite control network. *Journal of Scheduling*, 7:7–34.

[2] Bell, C. E. (1993). Scheduling deep-space network data transmissions: a Lagrangian relaxation approach. In *Proc. of SPIE 1963, Applications of Artificial Intelligence*, pages 330–340.

[3] Chien, S., Lam, R., and Vu, Q. (1997). Resource scheduling for a network of communication antennas. In *Proc. of the IEEE Aerospace Conference*, pages 361–373.

[4] Chien, S., Rabideau, G., Knight, R., R.Sherwood, Engelhardt, B., D.Mutz, T.Estlin, B.Smith, F.Fisher, T.Barrett, G.Stebbins, and D.Tran (2000). ASPEN: Automated planning and scheduling for space mission operations. In *Proc. of SpaceOps-00*.

[5] Damiani, S., Dreihahn, H., Noll, J., Niézette, M., and Calzolari, G. P. (2006). Automated allocation of ESA ground station network services. In *Proc. of IWPSS-06*.

[6] Gooley, T. (1993). Automating the satellite range scheduling process. Master's thesis, Air Force Institute of Technology.

[7] Guillaume, A., Seugnwon, L., Wang, Y.-F., Zheng, H., Hovden, R., Chau, S., Tung, Y.-W., and Terrile, R. (2007). Deep space network scheduling using evolutionary computational methods. In *Proc. of IEEE Aerospace Conference*, pages 1–6.

[8] Hoffmann, A., Dreihahn, H., Niézette, M., and Theis, G. (2009). Improving performance and interoperability of the estrack planning system. In *Proc. of IWPSS-09*.

[9] Johnston, M. and Clement, B. J. (2005). Automating deep space network scheduling and conflict resolution. In *Proc. of i-SAIRAS-05*.

[10] Johnston, M. and Tran, D. (2011). Automated scheduling for NASA's deep space network. In *Proc. of IWPSS-11*.

[11] Johnston, M., Tran, D., Arroyo, B., and Page, C. (2009). Request-driven scheduling for NASA's deep space network. In *Proc. of IWPSS-09*.

[12] Lenzen, C., Wörle, M. T., Mrowka, F., and Geyer, M. P. (2011). Automated scheduling for the TerraSar-X/Tandem-X. In *Proc. of IWPSS-11*.

[13] Parish, D. (1994). A genetic algorithm approach to automating satellite range scheduling. Master's thesis, Air Force Institute of Technology.

[14] Schalck, S. (1993). Automating satellite range scheduling. Master's thesis, Air Force Institute of Technology.

[15] Schmidt, M. and Schilling, K. (2009). A scheduling system with redundant scheduling capabilities. In *Proc. of IWPSS-09*.

[16] Schmidt, M. and Schilling, K. (2011). Enhanced redundant scheduling capability for low cost ground station networks. In *Proc. of IWPSS-11*.

[17] Verfaillie, G., Pralet, C., Olive, X., Rainjonneau, S., and Sebbag, I. (2012). Planning acquisitions for an ocean global surveillance mission. In *Proc. of i-SAIRAS-12*.