

Constraint programming for planning test campaigns of communications satellites

Emmanuel Hébrard, Marie-José Huguet, Daniel Veysseire, Ludivine Sauvan,
Bertrand Cabon

► **To cite this version:**

Emmanuel Hébrard, Marie-José Huguet, Daniel Veysseire, Ludivine Sauvan, Bertrand Cabon. Constraint programming for planning test campaigns of communications satellites. *Constraints*, Springer Verlag, 2017, 22 (1), pp.73 - 89. <10.1007/s10601-016-9254-x>. <hal-01443817>

HAL Id: hal-01443817

<https://hal.laas.fr/hal-01443817>

Submitted on 23 Jan 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Constraint Programming for Planning Test Campaigns of Communications Satellites

Emmanuel Hebrard¹, Marie-José Huguet¹, Daniel Veysseire¹, Ludivine Boche Sauvan^{1,2}, and Bertrand Cabon²

¹LAAS-CNRS, Université de Toulouse, CNRS, INSA, Toulouse, France

²Airbus Defence & Space, Toulouse, France

Abstract. The payload of communications satellites must go through a series of tests to assert their ability to survive in space. Each test involves some equipment of the payload to be active, which has an impact on the temperature of the payload. Sequencing these tests in a way that ensures the thermal stability of the payload and minimizes the overall duration of the test campaign is a very important objective for satellite manufacturers.

The problem can be decomposed in two sub-problems corresponding to two objectives: First, the number of distinct configurations necessary to run the tests must be minimized. This can be modeled as packing the tests into configurations, and we introduce a set of implied constraints to improve the lower bound of the model.

Second, tests must be sequenced so that the number of times an equipment unit has to be switched on or off is minimized. We model this aspect using the constraint SWITCH, where a buffer with limited capacity represents the currently active equipment units, and we introduce an improvement of the propagation algorithm for this constraint.

We then introduce a search strategy in which we sequentially solve the sub-problems (packing and sequencing). Experiments conducted on real and random instances show the respective interest of our contributions.

1 Introduction

The payload of a communications satellite is the on-board equipment that is actually relevant to the mission: receiving, amplifying and returning the signal. A set of tests are necessary to certify that the payload will correctly perform its mission. A test is characterized by a set of active equipment units and several thermal constraints limit the number of equipment units that can be made active simultaneously. The duration of the tests themselves is incompressible. However, activating some equipment takes time as each equipment unit must reach a given temperature and the temperature of the entire payload must be stabilized before tests can resume. Therefore, the total transition time depends on the order in which tests are sequenced. The goal is to sequence all the tests so that the overall duration is minimized.

Two main approaches have been previously considered. In the first approach [8], tests requiring the same subset of active equipment are packed together in the

same payload configuration. The transition time between tests run in the same configuration is null since no equipment activation is required. Between two configurations, however, some equipment units must be activated or deactivated and the heat of the payload must be stabilized before the next configuration, which entails important transition times. The objective of this approach is to minimize the number of configurations necessary for running all the tests. Moreover, a secondary objective is to minimize the overall number of activations and deactivations since the time required to stabilize the temperature of the payload is correlated, though not linearly, to the number of simultaneous activations. In practice, the transition time between two configurations is considered as a constant value sufficient to stabilize the temperature in the payload. Maillet et al. [8] first proposed a constraint programming approach to address this problem, using backjumping and dedicated heuristics.

It is in principle possible to activate an equipment unit whilst tests on other units are being run. The second approach, presented in [2], relies on this idea. An equipment activation is viewed as a task with a given duration. If an equipment unit is not tested during a period equal to this duration and if the thermal constraints allow it, then it can be activated at the beginning of this period, and becomes available for other tests at the end. It may be possible to do so for several or even all equipment activations, thus effectively masking the transition times. A local search method (simulated annealing) was proposed in [2] for this second approach, that is, with “online” activation allowed. It was shown that this approach can reduce the overall duration of the test campaign. However, the total number of activations can be higher in some instances.

Finally, since the focus in [8] was on selecting a subset of tests to be run rather than sequencing them, a straightforward improvement of the first approach was proposed in [2]: Once every test is allocated to a configuration and the number of configurations is minimized, the problem of minimizing the number of activations between consecutive configurations can be seen as a Traveling Salesman Problem (TSP). As the number of configurations is typically small, this problem can be solved optimally even with relatively basic TSP methods.

The second approach is difficult to implement in practice as activations and deactivations can happen continuously, making the detailed thermal analysis of the process a difficult task, whereas thermal engineers only need to worry about transitions between configurations in the first approach. In this work, we present a study for Airbus, which currently only implements the first approach. We propose an improved constraint programming approach to solve the problem.

Since thermal constraints limit the number of tests that can be run in the same configuration, this problem has a packing component, where tests have to be allocated to a minimum number of configurations. We introduce a set of implied constraints to improve the lower bound on this objective.

Moreover, this packing component is intertwined with a sequencing component where configurations must be ordered so that the number of equipment activations is minimized. We observe that this secondary objective can be easily modeled using the constraint SWITCH [1]. This constraint models a resource

operating as a buffer of limited capacity and whose content can be changed, however at a cost. Tests require equipment to be active and thermal constraints limit the number of simultaneously active equipment units, which can thus be seen as buffered. Moreover, the number of activations, which correspond to switches in the buffer, should be minimized. We introduce a simple improvement of the propagation algorithm for this constraint in the very common case where we have prior knowledge about the items that must be eventually buffered.

Next, we introduce a search strategy in which we sequentially solve the two sub-problems. We solve first the packing problem with a dedicated branching heuristic to find upper and lower bounds on this objective quickly. Then we solve a much simplified sequencing problem since tests are already allocated to configurations. This approach is not complete, and hence we must solve the overall problem in order to find optimal solutions. However this is made significantly easier thanks to the upper and lower bounds found in the previous phases.

Finally, we experimentally evaluate the different contributions and assess the benefit of our method with respect to the current approach in use at Airbus.

2 Formal Background

A constraint satisfaction problem (CSP) consists of a set of variables, where each variable X_k has a finite domain of values $D(X_k)$, and a set of constraints specifying allowed combinations of values for subsets of variables. A solution of a CSP is an assignment of values to the variables satisfying the constraints.

We consider both *integer* and *set* variables. A set variable Y_i is represented by its lower bound \underline{Y}_i which contains the required elements and an upper bound \overline{Y}_i which contains the possible elements. For a finite universe $U \subset \mathbb{N}$, we identify a set variable Y_i with the set of Boolean variables $\{Y_i^j \mid j \in U\}$. The predicates $Y_i^j = 1$ and $j \in Y_i$ are equivalent, as are $|Y_i| = \kappa$ and $\sum_{j \in U} Y_i^j = \kappa$.

For two integers $a \leq b$, we denote $[a, b]$ the set of consecutive integers $\{a, \dots, b\}$, and use the shortcut notation $[b]$ for $[1, b]$.

We shall often denote a sequence of variables or constants (c_1, \dots, c_n) by c (where the length n is either recalled or clear from the context).

2.1 The constraints SWITCH and BUFFEREDRESOURCE

The constraints SWITCH and BUFFEREDRESOURCE [1] were introduced to model a type of resource corresponding to a *buffer* which must contain the *items* required by some tasks when they are being processed. Such resources are limited in two ways: first, the buffer can only hold a limited number of items, and second, there is an upper bound on the number of item *switches* along the sequence.

In our context, these constraints are useful to model thermal constraints together with the objective to minimize the number of activations (switches) of equipment units (items to be buffered). The constraint SWITCH involves a sequence of set variables $Y = (Y_1, \dots, Y_n)$ and an integer variable M . The set variable Y_i represents the content of the buffer at position (or time) i , and the

variable M represents the total number of items that are removed from the buffer to make room for new items: $\sum_{i=1}^{n-1} |Y_i \setminus Y_{i+1}|$. Moreover, the buffer has a minimum and maximum capacity (lower and upper bound on $|Y_i|$) which are allowed to be different at every position i . Let Y be a sequence of set variables, and $\underline{\kappa}, \bar{\kappa}$ be two sequences of constants of same size n .

Definition 1 (SWITCH).

$$\text{SWITCH}(Y, \underline{\kappa}, \bar{\kappa}, M) \iff \forall i \in [n], \underline{\kappa}_i \leq |Y_i| \leq \bar{\kappa}_i \wedge \sum_{1 \leq i < n} |Y_{i+1} \setminus Y_i| \leq M$$

Often, we know beforehand which tasks are to be performed and which items are required by each task. In this case, one can use the `BUFFEREDRESOURCE` constraint which involves a sequence of integer variables $X = (X_1, \dots, X_n)$ representing a permutation of n tasks, and a sequence $t = (t_1, \dots, t_n)$ of sets of integers standing for the items required by each task. Achieving arc consistency on this constraint is NP-hard, and there is no dedicated propagation algorithm for this constraint, besides the obvious decomposition using `SWITCH`, `ALLDIFFERENT` [12] and some `ELEMENT` [7] constraints.

Definition 2 (BUFFEREDRESOURCE).

$$\begin{aligned} \text{BUFFEREDRESOURCE}(X, Y, \underline{\kappa}, \bar{\kappa}, t, M) \iff \\ & \text{SWITCH}(Y, \underline{\kappa}, \bar{\kappa}, M) \wedge \\ & \forall i < j \in [1, h], X_i \neq X_j \wedge \\ & \forall i, t_i \subseteq Y_{X_i} \end{aligned}$$

In our test planning problem, the buffers have equal upper and lower bounds. We shall therefore use a single integer parameter to denote the sequences $\underline{\kappa}$ and $\bar{\kappa}$ in the remainder of the paper.

3 Test Planning

3.1 Data and Constraints

A test campaign involves n tests and m equipment units. Every test $k \in [n]$ involves a subset $t_k \subseteq [m]$ of equipment units to be active.¹ A payload configuration (or simply configuration) is defined by a partition of the equipment into active and inactive units. A test k can occur in a configuration if the set of active equipment units in that configuration is a superset of t_k .

However, one cannot use a single configuration where every equipment unit is active, because the payload would overheat. Equipment units that are on the same wall or blade of the satellite contribute to the overall temperature of that wall/blade. Therefore, we have p constraints, one for every set of equipment units

¹ Throughout the paper, an equipment unit is said to be “active” if it is switched on and “inactive” otherwise.

whose thermal profiles are linked. For each thermal constraint $\ell \in [p]$, we define a subset $c_\ell \subseteq [m]$ of size Δ_ℓ of equipment units, among which exactly² κ_ℓ should be active at the same time, i.e., in the same configuration.

3.2 Decisions and Objectives

A test plan with h^* configurations is a mapping τ from tests to a set of consecutive integers $[h^*]$ (without loss of generality, we assume that configurations are numbered 1 to h^*), and a mapping σ from configurations to subsets of equipment units such that the equipment units required to run a test are active when this test is run and every configuration satisfies all thermal constraints.

The main objective is to minimize the number of configurations $|\{\tau(k) \mid k \in [n]\}|$, in order to reduce the transition time between tests, that is, the time spent in reconfiguring the payload.

Moreover it is important to take into account the total number of changes in the status of an equipment unit. Indeed, even though several equipment units can be switched on or off simultaneously, changing the status of more units requires a more careful analysis of the thermal dynamics of the system and is more likely to destabilize it. The second objective therefore is the total number of times an equipment unit is switched on besides the initial activation: $\sum_{i=1}^{h^*} (|\sigma(i) \setminus \sigma(i-1)|) - m$, where $\sigma(0)$ is assumed to be empty.

Indeed, it is important to take into account the total number of changes in the status of the equipment. Even though several equipment units can be switched on or off simultaneously, changing the status of more units requires a more careful analysis of the thermal dynamics of the system and is more likely to destabilize it. In order to count the total number of times each equipment unit is switched from inactive to active and vice versa, we must decide the order in which the planned configurations will be visited. By convention, since configuration names are arbitrary, the tests allocated to configuration i are run at the i -th position. Therefore, the same mapping τ defines both the allocation of tests to configurations and the sequence in which tests will be run.

We consider here that the packing objective has higher priority than the sequencing objective and thus that they are lexicographically ordered.

Example 1. Consider a set of 8 tests and 6 equipment units shown in Figure 1a. The equipment required by each test is indicated with the symbol \blackstar . Moreover, assume that we have two thermal constraints with scopes $c_1 = \{1, 2, 3\}$ and $c_2 = \{4, 5, 6\}$ both of capacity 2.

The solution shown in Figure 1a (\times symbols indicate active equipment not involved in the current test) is suboptimal as it requires three configurations. Additionally, equipment units 1 and 6 must be activated twice.

However, with the permutation 2, 3, 5, 7, 8, 1, 4, 6 shown in Figure 1b, we only need two configurations and every equipment unit is activated exactly once.

² Alternatively, one may only consider an upper bound only to prevent the system from overheating, however thermal engineers advise to keep the system as stable as possible, hence our choice of an equality.

| | | conf: 1 | | 2 | | | 3 | | |
|-----------|---|-----------|---|---|---|---|---|---|---|
| | | test: 1 2 | | 3 | 4 | 5 | 6 | 7 | 8 |
| equipment | 1 | ✖ | ✖ | | | | ✖ | ✖ | ✖ |
| | 2 | ✖ | ✖ | ✖ | ✖ | ✖ | | | |
| | 3 | | | ✖ | ✖ | ✖ | ✖ | ✖ | ✖ |
| | 4 | ✖ | ✖ | ✖ | ✖ | ✖ | | | |
| | 5 | | | ✖ | ✖ | ✖ | ✖ | ✖ | ✖ |
| | 6 | ✖ | ✖ | | | | ✖ | ✖ | ✖ |

(a)

| | | conf: 1 | | | | | 2 | | |
|-----------|---|---------|---|---|---|---|---|---|---|
| | | 2 | 3 | 5 | 7 | 8 | 1 | 4 | 6 |
| equipment | 1 | | | | | | ✖ | ✖ | ✖ |
| | 2 | ✖ | ✖ | ✖ | ✖ | ✖ | | | |
| | 3 | ✖ | ✖ | ✖ | ✖ | ✖ | ✖ | ✖ | ✖ |
| | 4 | | | | | | ✖ | ✖ | ✖ |
| | 5 | ✖ | ✖ | ✖ | ✖ | ✖ | ✖ | ✖ | ✖ |
| | 6 | ✖ | ✖ | ✖ | ✖ | ✖ | | | |

(b)

Fig. 1: A solution with 3 configurations and 2 extra activations (1a), and an optimal solution with 2 configurations and no extra activation (1b).

3.3 Complexity

It is relatively easy to see that the test planning problem described above is NP-hard. We show that the decision version TESTPLANNING, which asks whether there exists a plan with at most h configurations is NP-complete.

Theorem 1. TESTPLANNING is NP-complete.

Proof. It is in NP since a plan can be checked in polynomial time.

To prove hardness, we use a straightforward reduction from 3-COLORING, which asks, given a graph $G = (V, E)$, whether there exists a coloring of V with at most 3 colors such that no edge has its two end points of the same color.

From a graph $G = (V, E)$, we build an instance of TESTPLANNING as follows:

- For every edge $(x, y) \in E$, there are two equipment units x_y and y_x and a thermal constraint on these two units with capacity 1.
- For every vertex $x \in V$ we create a test t_x involving equipment units x_y for every y such that $(x, y) \in E$.

It is easy to see that two tests t_x and t_y can share the same configuration if and only if there is no edge $(x, y) \in E$. Therefore, G has a 3-coloring if and only if there is a plan with at most 3 configurations, and hence TESTPLANNING is NP-hard for $h = 3$. \square

Moreover, even if we let the number of configurations free, minimizing the number of switches is also NP-hard for a single thermal constraint over all equipment of capacity κ since it corresponds exactly to the constraint:

$$\text{BUFFEREDRESOURCE}(X, Y, \kappa, t, M)$$

where X (resp. Y) is a sequence of n integer (resp. set) variables, and for every $k \in [n]$ the variable X_k has domain $[m]$ and the variable Y_k ranges between the empty set and $[m]$.

The number of simultaneously active equipment units must be equal to κ_ℓ , therefore, we can consider actived equipment as a buffer of capacity exactly κ_ℓ . Moreover, the set of equipment units t_k required by a test k corresponds to the items required to be in the buffer when processing this task. Finally, the objective is to minimize the number activations which is equivalent, up to a constant, to the number of switches M . The reduction from Hamiltonian path to BUFFEREDRESOURCE provided in [1] can be lifted to this particular case.

4 A Constraint Model for Test Planning

Let $h \leq n$ be some known upper bound on the number of configurations. We use n *allocation* variables $\{X_k \mid k \in [n]\}$ with domain $[h]$ standing for the configuration allocated to each test. Unless we have a valid upper bound, the maximum number of configurations h is equal to n . Next, we use $m \times h$ *activity* variables $\{Y_i^j \mid j \in [m], i \in [h]\}$ standing for the status of equipment unit j in configuration i , i.e., Y_i^j is equal to 1 if unit j in switched ON in i and is equal to 0 otherwise. Moreover, we denote Y_i^U the set variable with characteristic function $\{Y_i^j \mid j \in U\}$. For instance, $Y_i^{[m]}$ has as characteristic function the set of variables in the column i of the $m \times h$ matrix formed by the activity variables, i.e., the set of equipment units active in configuration i . We introduce a variable M_ℓ for each constraint ℓ standing for the total number of switches on the equipment units c_ℓ . Finally, we have two variables to express the objective function: a variable N standing for the number of configurations and a variable M for the total number of switches.

To model the test planning problem, we then post the following constraints:

$$\forall k \in [n], \quad t_k \subseteq Y_{X_k}^{[m]} \quad (1)$$

$$\forall \ell \in [p], \forall i \in [h], \quad |Y_i^{c_\ell}| = \kappa_\ell \quad (2)$$

$$\forall k \in [n], \quad N \geq X_k \quad (3)$$

$$\forall \ell \in [p], \quad \text{SWITCH}(Y^{c_\ell}, \kappa_\ell, M_\ell) + \kappa_\ell - \Delta_\ell \quad (4)$$

$$M = \sum_{\ell=1}^p M_\ell \quad (5)$$

Constraints (1) channel the allocation variables and the activity variables to ensure that every equipment unit required by a test is active in the slot in which the test is ran. They are implemented as ELEMENT constraints.

Constraints (2) enforce the thermal constraints on the number of equipment units active simultaneously within given subsets. They are implemented as simple SUM constraints.

Constraints (3) ensure that the variable representing the number of configurations is greater than or equal than the maximum allocated slot. We do not use a MAXIMUM constraint nor an equality as we minimize this criterion.

Constraints (4) count the number of switches in the sequence of set variables $Y^{c_\ell} = (Y_1^{c_\ell}, \dots, Y_h^{c_\ell})$ for every thermal constraint ℓ . The SWITCH constraint standing for the thermal constraint ℓ ensures that the number of switches

$\sum_{i=1}^h |Y_i^{c_\ell} \setminus Y_{i+1}^{c_\ell}| = M_\ell$, for a capacity κ_ℓ of the buffer. The constant term $\kappa_\ell - \Delta_\ell$ is used to count only from the *second* activation of each equipment unit. The constraint SWITCH counts the number of deactivations. Moreover, the κ_ℓ items contained in the buffer at the last position are not counted by SWITCH, even though they will eventually be switched off. The total number of deactivations is thus $M_\ell + \kappa_\ell$ and it is equal to the number of activations. Then each of the Δ_ℓ equipment units constrained by ℓ must be activated at least once, so we can subtract this number to obtain the number of activation besides the first.

Finally, Constraint (5) computes the overall sum of switches.

We consider the number of configurations as a higher priority objective than the total number of switches. Therefore, we express the objective function to minimize as the weighted sum $(mh/2)N + M$.

The Test Planning problem may be decomposed into two sub-problems:

- a *Test Packing problem* when restricting to Constraints (1), (2) and (3) and the minimization of the number of configurations N ;
- a *Test Sequencing problem* when restricting to Constraints (1), (4) and (5) and the minimization of the total number of switches M .

For the type of satellites we have considered in this study, there is no overlap in the scope of thermal constraints since they stand for the physical support of a disjoint subset of equipment units. However, overlaps may exist in some satellite architectures. In this case, we need to replace the constraints (4) and (5) by a single SWITCH constraint on the sequence of set variables $Y^m = (Y_1^{[m]}, \dots, Y_h^{[m]})$:

$$\text{SWITCH}(Y^{[m]}, \sum_{\ell=1}^p \kappa_\ell, M) \tag{6}$$

5 Lower Bound for Test Packing

The problem restricted to Constraints (1), (2) and (3) has some similarities with the multi-dimensional bin packing problem [5]. Instead of a one dimensional capacity, the bins have a capacity κ_ℓ for each of the p dimension/thermal constraint. Moreover, for a test k requiring a set of equipment units t_k , we can compute a p -vector representing its weight in each of these dimensions. However, the weights are not additive since every equipment unit is activated at most once per configuration, irrespective of the number of tests requiring it in this configuration. As discussed previously, it can also be seen as a coloring problem. The particular case where each test is on a single equipment unit could be seen as a generalization of list coloring to hypergraphs, as thermal constraints can be mapped to hyperedges. However, to our knowledge, there is no known method for this specific problem.

We therefore use a simple lower bound on the number of configurations based on the “capacity” of the thermal constraints. Indeed, if a thermal constraint ℓ

ensures that at most κ_ℓ equipment units are active in a given configuration, and if the tests involve Δ_ℓ units, then at least $\lceil \Delta_\ell / \kappa_\ell \rceil$ configurations are needed. More generally, suppose that we have a lower bound K_j on the number of configurations in which an equipment unit j will be active. Given a thermal constraint ℓ , the sum of these lower bounds over all equipment units in c_ℓ cannot be greater than the number of configurations multiplied by the capacity κ_ℓ , in other words, for any constraint ℓ , $\lceil (\sum_{j \in c_\ell} K_j) / \kappa_\ell \rceil$ is a valid lower bound for N .

Now, consider an equipment unit j , and let $\Gamma_j = \{j' \mid \exists k \text{ s.t. } \{j, j'\} \subseteq t_k\}$ be its “neighborhood”, that is, the set of units that will necessarily be active when running the tests requiring j . It might not be possible to activate all these equipment units within a single configuration because of the thermal constraints. More generally, the equipment unit j must be active in at least as many configurations as are required to visit Γ_j , i.e., $K_j \geq \max(\lceil \frac{|\Gamma_j \cap c_\ell|}{\kappa_\ell} \rceil \mid \ell \in [p])$, let $\underline{\Gamma}_j$ be this lower bound.

The following two implied constraints enforce this lower bound:

$$\forall j \in [m], \quad K_j = \max(\underline{\Gamma}_j, \sum_{i=1}^h Y_i^j) \quad (7)$$

$$\forall \ell \in [p], \quad N \geq \left\lceil \frac{\sum_{j \in c_\ell} K_j}{\kappa_\ell} \right\rceil \quad (8)$$

However, Constraints 2 in the base model force every set variable $Y_i^{c_\ell}$ to have cardinality κ_ℓ , even if no test is allocated to configuration i . In this case, the configuration i would be a copy of configuration N since it satisfies every constraint and minimize the number of switches. However, this is not compatible with Constraint 7 since equipment units active in “dummy” configurations would still be counted. We therefore replace Constraints 2 with the following constraints on ph extra Boolean variables $\{Z_i^\ell \mid i \in [h], \ell \in [p]\}$:

$$\forall \ell \in [p], \forall i \in [h], \quad Z_i^\ell \iff \sum_{j \in c_\ell} Y_i^j > 1 \quad (9)$$

$$\forall \ell \in [p], \forall i \in [h], \quad \sum_{j \in c_\ell} Y_i^j = Z_i^\ell \times \kappa_\ell \quad (10)$$

$$\forall \ell \in [p], \forall i \in [h], \quad N \geq i \iff Z_i^\ell \quad (11)$$

Constraints 9 ensure that Z_i^ℓ equals 1 iff at least one equipment unit constrained by ℓ is active in configuration i . Constraints 10 ensure that, for every configuration i , either there is no active equipment (i is a dummy configuration) or exactly κ_ℓ for every thermal constraint ℓ . Finally, Constraints 11 channel these extra variables with the objective variable N by ensuring that $(Z_1^\ell, \dots, Z_h^\ell)$ is its unary *order* encoding [13]. These constraints are necessary to obtain a filtering as strong as in the base model, while improving the lower bound on N .

6 Test Sequencing

The problem of minimizing the number of equipment activations can be naturally represented using the SWITCH constraint, since, as shown in Section 4, thermal

constraints and the fact that tests require some equipment units to be active can be modeled as buffered resources. Moreover, as often in problems that can be represented with a buffered resource, we know beforehand the set of items (here equipment units) that will be activated at least once. In other words, the total number of items to be buffered minus the capacity κ_ℓ of the buffer is a trivial lower bound on the required number of switches. However, the constraint SWITCH does not take this information into account and hence is “suboptimal”, especially when only a few tests have been allocated a configuration.

In this section, we propose an improvement of the propagator for SWITCH for the very common case where we have prior knowledge on a set of items that must eventually be buffered. In some cases, we can simply add the number of non-buffered items to the lower bound computed by the algorithm for SWITCH. However, this is not always true. We define a correct lower bound based on this idea in Theorem 2.

We define the variant SWITCH⁺ of the SWITCH constraint with an extra parameter V indicating the items that must be put in the buffer at some point in the sequence. Let V be a set of integers, M an integer variable, $Y = (Y_1, \dots, Y_n)$ be a sequence of set variables, and $\underline{\kappa}, \bar{\kappa}$ be two sequences of n integers.

Definition 3.

$$\text{SWITCH}^+(Y, \underline{\kappa}, \bar{\kappa}, V, M) \iff \forall i \in [n], \underline{\kappa}_i \leq |Y_i| \leq \bar{\kappa}_i \wedge \sum_{1 \leq i < n} |Y_{i+1} \setminus Y_i| \leq M \wedge \bigcup_{i=1}^n Y_i = V$$

We next recall some background about the constraint and in particular its propagation algorithm in Definitions 4 and 5. It is possible to find an optimal buffer sequence σ , that is, an assignment of Y that minimizes the value of M with the greedy procedure **FindSupport** introduced by [1] (Algorithm 1 in [1]). This algorithm explores the sequence once while maintaining all items in a list ordered by a priority relation \prec based on two indices ($\text{next}_\in^i(j)$ and $\text{next}_\notin^i(j)$) for each item j .

Definition 4. $\text{next}_\in^i(j)$ is the least index $i' \geq i$ such that $j \in \underline{Y}_{i'}$ if it exists and $n + 2$ otherwise. $\text{next}_\notin^i(j)$ is the least index $i' \geq i$ such that $j \notin \bar{Y}_{i'}$ if it exists and $n + 1$ otherwise.

The priority relation \prec between two items is defined by the following criteria:

Definition 5. At index i , and given two items $j_1 < j_2$, we have $j_1 \prec j_2$ if:

1. $\text{next}_\in^i(j_1) < \text{next}_\notin^i(j_1)$ and $\text{next}_\in^i(j_1) \leq \text{next}_\in^i(j_2)$, or
2. $\text{next}_\in^i(j_2) > \text{next}_\notin^i(j_2)$ and $\text{next}_\notin^i(j_1) \geq \text{next}_\notin^i(j_2)$

and $j_2 \prec j_1$ otherwise.

The procedure starts with $\sigma(0) = \emptyset$ and when moving to index i , it first adds all required items (i.e., in \underline{Y}_i) and removes all impossible items (i.e., not in \bar{Y}_i),

which yields the set $\sigma(i) = \bar{Y}_i \cap \sigma(i-1) \cup \underline{Y}_i$. If $|\sigma(i)| > \bar{\kappa}_i$ it then removes the $\bar{\kappa}_i - |\sigma(i)|$ last items for the order \prec in $\sigma(i)$. Otherwise, if $|\sigma(i)| < \underline{\kappa}_i$ it adds the $\underline{\kappa}_i - |\sigma(i)|$ first items for the order \prec in $\bar{Y}_i \setminus \sigma(i)$.

Definition 6. A stretch of a buffer sequence σ is a triple $\langle j, a, b \rangle$ such that the value j is buffered in σ during the interval $[a, b]$ and j is not buffered at $a - 1$ nor at $b + 1$ (we assume that no item is buffered at 0 or $n + 1$). We say that a stretch $\langle j, a, b \rangle$ is optional if there is no $i \in [a, b]$ such that $j \in \underline{Y}_i$. We say that the item j is optional if there exists a stretch $\langle j, a, b \rangle$ and for all $i \in [n]$, $j \notin \underline{Y}_i$.

Every stretch entails one switch, except if it extends until the end of the sequence. Therefore, if κ is the cardinality of the buffer at the last index of the sequence, the following observation holds:

Observation 1 A solution with α switches has $\kappa + \alpha$ stretches.

Lemma 1. If σ and σ' are two buffer sequences found by **FindSupport** on two instances I and I' equal on every index except one for which the lower bound of the buffer at this index contains a single additional item j in I' , then $\forall i \in [n]$, $\sigma'(i) \setminus \sigma(i) \subseteq \{j\}$.

Proof. An item not in the buffer at index $i - 1$ is added at index i only if:

1. the item is the lower bound \underline{Y}_i or
2. $\underline{\kappa}_i > |\bar{Y}_i \cap \sigma(i-1) \cup \underline{Y}_i|$ and it is in the $|\bar{Y}_i \cap \sigma(i-1) \cup \underline{Y}_i| - \bar{\kappa}_i$ first for \prec .

Only item j can satisfies case 1 in I' but not in I .

Moreover, by definition, the order \prec is equal in I and I' , except for j which may be ranked higher in I' . Therefore, the only item that may satisfy case 2 in I' but not in I is again j . Therefore, $\forall i \in [n]$, $\sigma'(i) \setminus \sigma(i) \subseteq \{j\}$. \square

Lemma 2. Let σ be a buffer sequence of an instance I with minimal number of switches and let j be an item not buffered in σ . For any $i \in [n]$, a sequence σ' on the instance I' obtained by adding the constraint $j \in \underline{Y}_i$ has at least one more non-optional stretch than σ .

Proof. We can assume that σ and σ' were found by **FindSupport** as this algorithm is complete. The sequence σ' has necessarily a non-optional stretch $\langle j, a, b \rangle$, and there was no j -stretch in σ . Therefore, if the number of non-optional stretches is not larger in σ' than in σ , it must have one less non-optional j' -stretch for an item $j' \neq j$. This can only happen if the gap between two non-optional stretches $\langle j', a_1, b_1 \rangle$ and $\langle j', a_2, b_2 \rangle$ with $b_1 + 1 < a_2$ is bridged by buffering the value j' in the interval $[b_1 + 1, a_2 - 1]$. Hence there exists i such that $j' \in \sigma'(i) \setminus \sigma(i)$. However, by Lemma 1 we have $\forall i \in [n]$, $\sigma'(i) \setminus \sigma(i) \subseteq \{j\}$. \square

Theorem 2. For any two sets $V \subseteq U$, if there is an optimal buffer sequence visiting exactly the items in $U \setminus V$ with α switches, β optional stretches, and γ optional items, then there is no buffer sequence visiting all items in U in less than $\alpha + |V| - \beta + \gamma$ switches.

Proof. First, notice that we can reduce the case with $\gamma > 0$ optional items to the case without optional item, and $\beta - \gamma$ optional stretches. Indeed, if there exists an optimal sequence σ visiting all items in U , we can add a constraint $j \in \underline{Y}_i$ for every pair (i, j) where $j \in U$ and $j \in \sigma(i)$. The procedure **FindSupport** will then find a sequence with same number of switches, same number of stretches, and γ less optional stretches than σ . Therefore, we suppose $\gamma = 0$ and prove the lower bound $\alpha + |V| - \beta$.

Now, if σ has $\kappa + \alpha - \beta$ non-optional stretches, then by Lemma 2, we know that a solution visiting all items in U must have at least $\kappa + \alpha + |V| - \beta$ (non-optional) stretches, and hence at least $\alpha + |V| - \beta$ switches. \square

Counting the number of optional stretches and items in the sequence returned by **FindSupport** can be done in linear time. Therefore, Theorem 2 improves the lower bound found by this algorithm without changing its worst case complexity when we know that some items must be buffered but do not appear in the lower bound of a set variable. This is true in the test sequencing problem, as it is in most applications of this constraint.

7 Search Strategy

In this section we introduce a dedicated branching heuristic for the packing problem. The basic idea is that we can easily evaluate the impact of allocating a test to a configuration by counting how many equipment activations are required and how these equipment units are already constrained in this configuration. Second, we propose to decompose the problem into packing and sequencing aspects in order to find good upper bounds quickly for the test planning problem.

7.1 Branching Heuristic

Let $\delta_k^\ell(i)$ be the number of equipment units constrained by ℓ that will be active in configuration i if test k was to be run in that configuration, i.e., the number of non-ground Boolean activation variables concerning equipment of test k constrained by ℓ :

$$\delta_k^\ell(i) = \sum_{j \in c_\ell \cap t_k} \bar{Y}_j^i - \underline{Y}_j^i$$

We consider the ratio $\frac{\Delta_\ell}{\kappa_\ell}$ to be proportional to the tightness of the constraint ℓ , and we use the change in tightness resulting from the decision of running test k in configuration i to evaluate the impact of this decision. After the decision $X_k = i$, the tightness $r_b = \frac{\Delta_\ell}{\kappa_\ell}$ becomes $r_a = \frac{\Delta_\ell - \delta_k^\ell(i)}{\kappa_\ell - \delta_k^\ell(i)}$. Therefore, the factor r_a/r_b represents the factor by which the tightness of constraint ℓ would increase. As $r_a/r_b \in [1, \infty]$, we use 1 minus its inverse as a measure, in $[0, 1]$, of the impact of that decision, that is:

$$\gamma_k^\ell(i) = 1 - \frac{\Delta_\ell(\kappa_\ell - \delta_k^\ell(i))}{\kappa_\ell(\Delta_\ell - \delta_k^\ell(i))}$$

Now, we can use the average of these impacts on all thermal constraints to define the impact $\gamma(X_k \leftarrow i)$ of allocating test k to configuration i .

$$\gamma(X_k \leftarrow i) = \frac{\sum_{\ell \in [p]} \gamma_k^\ell(i)}{p}$$

Notice that allocation variables may have many more values than necessary to satisfy the thermal constraints. Indeed, it is important, for a coloring heuristic, to branch only values in $[i+1]$ where i is the highest allocated value so far. Hence, given a test k we consider the intersection $\widehat{D}(k) = D(X_k) \cap [i+1]$ instead of its actual domain $D(X_k)$. We therefore select the variable X_k minimizing:

$$\frac{|\widehat{D}(k)|}{\sum_{i \in \widehat{D}(k)} \gamma(X_k \leftarrow i)}$$

And we branch on the value i minimizing $\gamma_k^\ell(i)$.

7.2 Multi-stage Approach

We use mh Boolean variables to represent the status of every equipment unit in every configuration. Moreover, without an upper bound on the number of configurations required to pack every test, we can only assume that n configurations (as many as there are tests) may be needed, i.e., $h = n$. However, in practice n is a gross overestimate of h . For instance, in one of the industrial instances that we considered, several hundreds of tests can be run in as few as three configurations.

Furthermore, since we consider two criteria in a hierarchical way, it makes sense to optimize a relaxation of the problem where only Constraints 1 to 3 are kept (the packing sub-problem), i.e., the model is complete with respect to the objective with highest priority. Observe that since the order of the bins does not matter in this case, configurations are symmetric. We therefore used lexicographic ordering constraints [4, 6] on the set variables Y_1, \dots, Y_h .

Last, we also considered the pure sequencing aspect of the problem. Given a packing, finding the optimal sequence for that packing can be modeled with a set of BUFFEREDRESOURCE constraints, one for each thermal constraint, all sharing the same permutation. Another way to understand this is that we can consider a packing solution using h configurations as a new instance with only h tests to sequence (it is unlikely that two such tests could share a single configuration given the thermal constraints) and consider Constraints 1, 4 and 5 of the complete model (the sequencing sub-problem). Solving this problem to optimality does not give us a lower bound on the total number of switches, however, it is much simpler and can often provide a good upper bound quickly.

We therefore implemented the following four-phase strategy:

1. We run a greedy descent on the packing problem to find an initial upper bound. The trivial heuristic that branch on the lexicographically least configuration for a test gives relatively good results, so we stop this phase at the first solution found, which is backtrack-free.

2. We run the packing model (initialized with the previous upper bound) for a given period of time, or until optimality is proven.
3. We run the sequencing model for a given period of time, or until optimality is proven (though in this case we cannot deduce a lower bound).
4. We run the complete model (packing & sequencing) (initialized with lower and/or upper bounds, accordingly) for the rest of the allocated time.

8 Experimental Evaluation

We tested the different approaches that we propose on industrial and generated instances. We have only six real instances, corresponding to three already launched communications satellites. The same tests are usually run in two types of thermal environments. The `HOT` and `COLD` test phases respectively simulate the periods where the satellite is facing the sun, or when it is in Earth’s shadow. Instances labeled `cold` are much more thermally constrained and thus typically require more configurations than those labeled `hot`. We diversified the pool of instances by randomly shuffling tests in order to produce 5 randomized variants of every instance. Moreover, we used random instances designed to be similar to the real cases, generated as follows: for a given number of tests $n \in \{30, 50, 80, 100, 200, 300\}$, we set the number of equipment units to $n/4$. The equipment of instances with 30 and 50 tests are equally partitioned into 3 thermal constraints. Each test requires 2 equipment units from two different thermal constraints. Other instances have 5 thermal constraints and each test requires 2 or 3 equipment units belonging to different thermal constraints. Then, for each generated instance, we consider two levels of tightness for thermal constraints to simulate the `HOT` ($\kappa_\ell/\Delta_\ell = 0.6$) and the `COLD` phases ($\kappa_\ell/\Delta_\ell = 0.4$).

We generated 5 variants of 12 classes of instances. Instance `XXX.YY` denote a set of `XXX` tests with `YY` giving the ratio κ_ℓ/Δ_ℓ of the thermal constraints. Instances `A`, `B` and `C` are industrial instances. All models have been implemented in Choco 3 [10] and ran on Intel Xeon E5 processors for a total of thirty minutes on every instance. We compared the five following approaches:

- **base** is the straightforward model with Constraints 1 to 5, using *Weighted Degree* [3] for variable selection and lexicographic branching. Notice that the strategy is set up to branch on all allocation variables before branching on activity variables. Other predefined heuristics in Choco were less efficient. It is important to note that due to the “coloring” aspect of the problem, branching on the lexicographically least value (color) in the domain is extremely important. However *Impact Based Search* [11] and *Activity Based Search* [9] cannot trivially be made to branch on the lexicographically least value and thus gave extremely poor results for that reason.
- **heuristic** is the same model as **base**, however with the dedicated heuristic described in Section 7.1 for variable ordering.
- **multi-stage** is the same model as **base**, however, using the four-phase strategy described in Section 7.2.

Table 1: Methods comparison. Average number of configurations and switches on random and industrial instances.

| instance | base | | heuristic | | multi-stage | | propagation | | full | |
|----------|---------|---------|-----------|----------|-------------|---------|-------------|----------|---------|----------|
| | #conf | #switch | #conf | #switch | #conf | #switch | #conf | #switch | #conf | #switch |
| 030-04 | 6.4 (5) | 1.6 (5) | 6.4 (5) | 1.6 (5) | 6.4 (5) | 1.6 (4) | 6.4 (5) | 1.6 (5) | 6.4 (5) | 1.6 (4) |
| 030-06 | 3.4 (5) | 0.8 (5) | 3.4 (5) | 0.8 (5) | 3.4 (5) | 0.8 (3) | 3.4 (5) | 0.8 (5) | 3.4 (5) | 0.8 (2) |
| 050-04 | 5.2 (2) | 2.2 (2) | 5.2 (3) | 2.0 (3) | 5.2 (5) | 2.0 (3) | 5.2 (5) | 2.0 (5) | 5.2 (5) | 2.0 (5) |
| 050-06 | 3.6 (2) | 0.8 (2) | 3.6 (5) | 0.6 (5) | 3.6 (5) | 0.6 (3) | 3.6 (5) | 0.6 (5) | 3.6 (5) | 0.6 (2) |
| 080-04 | 9.6 | 19.0 | 9.0 (1) | 15.0 (1) | 9.0 | 15.4 | 9.0 (1) | 13.8 (1) | 9.0 (3) | 13.8 (1) |
| 080-06 | 5.6 | 6.0 | 5.4 (3) | 4.8 (3) | 5.4 (3) | 5.2 (1) | 5.4 (3) | 4.2 (3) | 5.4 (4) | 4.2 (3) |
| 100-04 | 14.0 | 40.6 | 13.0 | 43.4 | 13.2 | 31.8 | 13.0 | 34.4 | 12.8 | 30.2 |
| 100-06 | 4.4 | 5.8 | 4.0 (4) | 2.8 (4) | 4.0 (5) | 4.0 | 4.0 (4) | 3.0 (4) | 4.0 (5) | 2.8 (5) |
| 200-04 | 12.6 | 67.2 | 10.2 | 62.8 | 10.2 | 46.2 | 10.4 | 53.6 | 10.0 | 47.8 |
| 200-06 | 5.8 | 24.2 | 5.0 | 18.2 | 5.0 | 17.2 | 5.0 | 14.0 | 5.0 | 14.4 |
| 300-04 | - | - | 10.4 | 111.2 | 10.2 | 85.0 | 10.2 | 94.0 | 10.0 | 81.8 |
| 300-06 | 5.7 | 45.3 | 4.8 | 37.0 | 4.6 | 26.8 | 4.2 | 25.4 | 4.0 | 22.4 |
| A cold | 6.0 | 16.0 | 6.0 | 14.2 | 6.0 | 8.8 | 6.0 | 9.6 | 6.0 | 7.0 |
| A hot | 4.0 | 2.0 | 4.0 | 4.6 | 4.0 | 2.0 | 4.0 | 2.0 | 4.0 (5) | 2.0 |
| B cold | - | - | 9.0 | 106.6 | 9.0 | 50.2 | 9.0 | 78.8 | 9.0 | 50.2 |
| B hot | - | - | 4.0 | 19.6 | 4.0 | 3.6 | 4.0 | 1.2 | 4.0 (5) | 0.8 (2) |
| C cold | - | - | 8.0 | 82.8 | 8.0 | 46.8 | 8.0 | 64.8 | 8.0 | 50.6 |
| C hot | - | - | 3.0 | 15.0 | 3.0 | 2.0 | 3.0 (4) | 0.2 (4) | 3.0 (5) | 0.0 (5) |

- **propagation** augments the model **base** with symmetry breaking, the lower bound on the packing defined by Constraints (7) to (11), and the improvement on the propagator for SWITCH described in Section 6.
- **full** is the same model as **propagation**, however it uses the four-phases strategy and the branching heuristic described in Section 7.

Note that for every approach, we first applied a preprocessing to the data in order to merge identical tests. Indeed, on top of the packing based on the configuration of active equipment units, one must, in the real setting, further partition the tests because of different requirements on signal routing equipment. This second packing phase is exactly a list-coloring problem on the tests of each the configurations. The size of these problems is very modest with respect to state of the art coloring algorithms, so we do not study this aspect in the current paper. As a consequence, some tests in the data require the same set of equipment units to be active, and can be thought of as a single test for our purpose.

The results are reported in Table 1. For each method, we show the objective values, where *#conf* stands for the average number of configurations over the five variants of the instance, and *#switch* stands for the average number of switches for the same instances. Next to these values, we report in brackets, the number of instances for which the reported objective value was proven optimal within a 30 minutes time limit (no value means that none of the runs was complete). For each instance, the methods proving optimality in the most cases, among those giving the minimum average objective value, are color-highlighted.

First, we observe that the straightforward model **base** is very poor, in most of the larger instances it does not even find a feasible solution in less than ten minutes. On the other hand, augmenting this model either with stronger

Table 2: Impact of the lower bounds. Average number of configurations and switches on random and industrial instances.

| instance | lb conf | | lb switch | | full \ lb switch | | full \ lb conf. | |
|----------|---------|----------|-----------|---------|------------------|----------|-----------------|----------|
| | #conf | #switch | #conf | #switch | #conf | #switch | #conf | #switch |
| 030-04 | 6.4 (5) | 1.6 (5) | 6.4 (5) | 1.6 (5) | 6.4 (5) | 1.6 (4) | 6.4 (5) | 1.6 (3) |
| 030-06 | 3.4 (5) | 0.8 (5) | 3.4 (5) | 0.8 (5) | 3.4 (5) | 0.8 (2) | 3.4 (5) | 0.8 (2) |
| 050-04 | 5.2 (3) | 2.2 (3) | 5.2 (4) | 2.0 (4) | 5.2 (5) | 2.0 (3) | 5.2 (5) | 2.0 (4) |
| 050-06 | 3.6 (4) | 0.6 (4) | 3.6 (5) | 0.6 (5) | 3.6 (5) | 0.6 (2) | 3.6 (5) | 0.6 (2) |
| 080-04 | 9.2 (1) | 16.6 (1) | 9.0 | 16.8 | 9.0 (3) | 14.4 (1) | 9.0 (3) | 13.4 (1) |
| 080-06 | 5.4 (3) | 5.2 (3) | 5.4 (3) | 4.4 (3) | 5.4 (4) | 4.6 (3) | 5.4 (5) | 3.8 (3) |
| 100-04 | 13.8 | 37.0 | 13.8 | 37.0 | 12.8 | 30.8 | 12.8 | 30.4 |
| 100-06 | 4.0 | 4.6 | 4.0 (1) | 4.4 (1) | 4.0 (5) | 2.8 (5) | 4.0 (5) | 2.8 (5) |
| 200-04 | 11.0 | 57.0 | 11.0 | 56.6 | 10.0 | 47.2 | 10.0 | 47.8 |
| 200-06 | 5.0 | 18.8 | 5.0 | 17.6 | 5.0 | 14.4 | 5.0 | 14.8 |
| 300-04 | 11.0 | 98.8 | 11.0 | 97.4 | 10.0 | 81.8 | 10.0 | 85.8 |
| 300-06 | 5.0 | 28.2 | 5.0 | 29.8 | 4.0 | 22.0 | 4.0 | 21.8 |
| A cold | 6.0 | 10.2 | 6.0 | 7.6 | 6.0 | 12.0 | 6.0 | 7.2 |
| A hot | 4.0 | 3.0 | 4.0 | 2.8 | 4.0 (5) | 3.2 | 4.0 (5) | 2.0 |
| B cold | 9.0 | 64.8 | 9.2 | 65.4 | 9.0 | 50.6 | 9.0 | 52.2 |
| B hot | 4.0 | 3.8 | 4.0 | 4.0 | 4.0 (5) | 1.8 | 4.0 | 2.4 |
| C cold | 8.0 | 55.4 | 8.0 | 55.2 | 8.0 | 49.2 | 8.0 | 50.0 |
| C hot | 3.0 | 5.8 | 3.0 | 5.2 | 3.0 (5) | 0.8 | 3.0 (5) | 0.4 (3) |

propagation, better branching heuristic, or the multi-stage approach is sufficient to obtain decent results. Combining all these improvements clearly yields the best and most robust results, and sometimes allows to prove optimality even on real industrial instances.

Second, multi-stage approaches (**multi-stage** and **full**) are clearly better for larger instances, however, they tend to hinder the ability of Choco to prove optimality on small instances. Moreover, we observed that they tend to be less robust, in the sense that the final result greatly depends on the initial packing which gives no guarantees on the number of switches. This is especially true for large instances in which the complete model often cannot improve on the heuristic sequence found during the third phase.

For a deeper analysis of the two bounds proposed in this paper, we ran four other models on the same instances. The first two are the base model augmented with the lower bound on configurations or the lower bound on switches (**base** \oplus **lb conf** and **base** \oplus **lb switch**, respectively). The other two models are the full model from which we removed these lower bounds (**full** \ **lb switch** and **full** \ **lb conf**, respectively). The results of these additional tests, in Table 2 clearly show that both bounds are useful. Surprisingly, the capacity to prove optimality on a criterion is also impacted by the bound on the other criterion.

Finally, in Table 3, we compare the results of our best method (model **full**) with the method previously used by Airbus Defense & Space [8] and with a slight improvement of this method described in [2]. The former, denoted **CM**, is a constraint optimization tool build to solve the packing problem only. The second, denoted **CM+TSP**, is the same approach, however the resulting configurations

Table 3: Comparison with current methods. Number of configurations and switches on industrial instances.

| instance | CM | | CM+TSP | | Choco (full) | |
|----------|-------|---------|--------|---------|--------------|---------|
| | #conf | #switch | #conf | #switch | #conf | #switch |
| A cold | 6 | 43 | 6 | 9 | 6 | 7.0 |
| A hot | 4 | 5 | 4 | 3 | 4 (5) | 2.0 |
| B cold | 10 | 108 | 10 | 70 | 9 | 50.2 |
| B hot | 5 | 30 | 5 | 1 | 4 (5) | 0.8 (2) |
| C cold | 8 | 91 | 8 | 41 | 8 | 50.6 |
| C hot | 3 | 14 | 3 | 9 | 3 (5) | 0.0 (5) |

are then permuted so that the *Traveling Salesman Problem* defined by the Hamming distance between configurations is optimized. It is not surprising that the first method is very poor in terms of total equipment activation since it makes no attempt to optimize this criterion. However, it is interesting to compare with our results as it is still the method used in practice.

The method **CM+TSP** gives a better solution in one case (**C cold**). Indeed, this instance is particular as it involves more than twice as many tests as other instances. The consequence is that the model computing packing and sequencing simultaneously is relatively inefficient. For this instance, we therefore ran a simple randomized sequence of the second and third phases of the multi-stage approach (i.e., pure packing followed by pure sequencing) and quickly found a similar solution (though we could not improve on it). Notice that in the case of the “hot” test campaign for satellite **C**, carefully packing and sequencing the tests makes it possible to get rid of all equipment activation besides the mandatory one, whereas 14 activations are necessary with the current method.

9 Conclusion

We have introduced a complete constraint programming approach for the problem of packing and sequencing the validation tests of communications satellites. We proposed a search strategy and lower bound for the packing aspect of the problem. Moreover, we introduced an improvement of the SWITCH constraint that can be applied in many other contexts. Our experimental evaluation shows that the methods proposed in this paper greatly improve the test plans with respect to those currently used within the Airbus group.

Although this approach is not yet industrially implemented, a previous internal study in Airbus³ has shown that, during a test campaign, around 30% of the total duration is spent on transitions between configurations. Moreover, in many cases, tests must be interrupted because the payload is overheating and

³ Master’s internship report by Ludivine Boche-Sauvan for the “Institut Supérieur de l’Aéronautique et de l’Espace” (ISAE) in 2012 (http://www.laas.fr/files/ROC/LAAS_Techreport.pdf).

can only resume after the system has been stabilized. The constraint model we introduced should help with both of these issues.

Since such test campaigns require an extremely costly and energy greedy thermal vacuum chamber as well as a large team of engineers in 3-shift rosters, significant financial savings are expected from this approach.

References

1. Christian Bessiere, Emmanuel Hebrard, Marc-André Ménard, Claude-Guy Quimper, and Toby Walsh. Buffered resource constraint: Algorithms and complexity. In *Proceedings of the Eleventh International Conference on Integration of AI and OR Techniques in Constraint Programming (CPAIOR)*, pages 318–333, 2014.
2. Ludivine Boche-Sauvan, Bertrand Cabon, Marie-José Huguet, and Emmanuel Hébrard. Heuristic Methods for Test Sequencing in Telecommunication Satellites. In *Proceedings of the Tenth International Conference on Modeling, Optimization and Simulation (MOSIM)*, 2014.
3. Frédéric Boussemart, Fred Hemery, Christophe Lecoutre, and Lakhdar Sais. Boosting Systematic Search by Weighting Constraints. In *Proceedings of the Sixteenth European Conference on Artificial Intelligence (ECAI)*, pages 146–150, 2004.
4. Mats Carlsson and Nicolas Beldiceanu. Arc-Consistency for a Chain of Lexicographic Ordering Constraints. Technical Report T2002-18, Swedish Institute of Computer Science, 2002.
5. Chandra Chekuri and Sanjeev Khanna. On multi-dimensional packing problems. In *Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms, Tenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 185–194, 1999.
6. Alan M. Frisch, Brahim Hnich, Zeynep Kiziltan, Ian Miguel, and Toby Walsh. Propagation algorithms for lexicographic ordering constraints. *Artificial Intelligence*, 170(10):803–834, 2006.
7. Pascal Van Hentenryck and Jean-Philippe Carillon. Generality versus Specificity: An Experience with AI and OR techniques. In *Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI)*, pages 660–664, 1988.
8. Caroline Maillet, Gérard Verfaillie, and Bertrand Cabon. Constraint programming for optimising satellite validation plans. In *Proceedings of Seventh International Workshop on Planning and Scheduling for Space (IWSPSS)*, 2011.
9. Laurent Michel and Pascal Van Hentenryck. *Activity-Based Search for Black-Box Constraint Programming Solvers*, pages 228–243. 2012.
10. Charles Prud’homme, Jean-Guillaume Fages, and Xavier Lorca. *Choco Documentation*. TASC, INRIA Rennes, LINA CNRS UMR 6241, COSLING S.A.S., 2015.
11. Philippe Refalo. Impact-Based Search Strategies for Constraint Programming. In *Principles and Practice of Constraint Programming (CP)*, pages 557–571, 2004.
12. Jean-Charles Régin. A filtering algorithm for constraints of difference in CSPs. In *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI)*, pages 362–367, 1994.
13. Naoyuki Tamura, Akiko Taga, Satoshi Kitagawa, and Mutsunori Banbara. Compiling finite linear CSP into SAT. *Constraints*, 14(2):254–272, 2008.