Computers & Operations Research 38 (2011) 1572–1583

Contents lists available at ScienceDirect

Computers & Operations Research

journal homepage: www.elsevier.com/locate/caor

A Lagrangian heuristic for satellite range scheduling with resource constraints

Fabrizio Marinelli [a,*], Salvatore Nocella [b], Fabrizio Rossi [b], Stefano Smriglio [b]

[a] Dipartimento di Ingegneria Informatica, Gestionale e Automazione Università Politecnica delle Marche Via Brecce Bianche, I-60131 Ancona, Italy
[b] Dipartimento di Informatica Università degli Studi di L'Aquila via Vetoio, I-67010 Coppito, L'Aquila, Italy

ARTICLE INFO

Available online 28 January 2011

Keywords:
Multiprocessor scheduling
Applications
Satellite communications
Time-indexed scheduling formulation
Lagrangian heuristic

ABSTRACT

The data exchange between ground stations and satellite constellations is becoming a challenging task, as more and more communication requests must be daily scheduled on a few, expensive stations located all around the Earth. Most of the scheduling procedures adopted in practice cannot cope with such complexity, and the development of optimization-based tools is strongly spurred.

We show that the problem can be formulated as a *multiprocessor task scheduling* problem in which each job (communication) requires a time dependent pair of resources (ground station and satellite) to be processed, and the objective consists of maximizing the total revenue of on-time jobs. A time-indexed {0,1}-linear programming formulation is then introduced able to include all the complex technological constraints of current constellations. Unfortunately, relevant real-world scenarios yield integer programs with hundreds of thousands variables and a few million constraints, which cannot be tackled by standard integer programming (either exact or heuristic) techniques.

To overcome this difficulty, we developed a Lagrangian version of the *Fix-and-Relax* MIP heuristic. It is based on a Lagrangian relaxation of the problem which is shown to be equivalent to a sequence of *maximum weighted independent set* problems on interval graphs. The heuristic has been implemented in a tool used by the Italian reference operator for the GALILEO constellation, providing near optimal solutions to relevant large scale test problems.

1. Introduction

Satellite systems support many different services, such as surveillance, geodesy and navigation, remote sensing and monitoring, telecommunications and data relay. The dramatic increase of the demand for such services is currently stimulating the implementation of new constellations. One major technological novelty is that the large satellites used in early (*monolithic*) constellations are replaced by clusters of smaller satellites within a single orbital location, yielding the so-called *distributed satellite constellations* [34]. A distributed constellation is *flexible*, as it simultaneously delivers a variety of services, and *robust*, as it provides a high level of service even when the constellation is not completely deployed or in case of some component failure. Moreover, small satellites are cheaper and can be managed more efficiently than large ones. On the other hand, the implementation of distributed satellite constellations creates challenging problems, particularly on the planning and control side. One major problem regards the data exchange between *ground stations* and spacecrafts. In fact, the set of communications (from now on referred to as *services*) to be scheduled everyday is rapidly getting larger and larger. On the contrary, the ground station networks consists of a few big and expensive installations, typically spread over different continents, and can hardly be expanded. In such a setting, careful scheduling policies are getting more and more important. Remarkably, when spacecraft health and orbit control services are considered (called *navigation* services), minimizing the number of tardy services is critical to keep safe both the constellation and the performance delivered. Unfortunately, these problems are quite difficult in practice. Indeed, the common procedures, in which a team of experts build schedules using worksheets or simple scheduling tools, often yield poor solutions or even fail to find a feasible one. This has been experienced by the engineers of Telespazio S.p.A.[1] while planning missions for the forthcoming GALILEO constellation.

* Corresponding author.
E-mail addresses: fabrizio.marinelli@univpm.it (F. Marinelli), nocella@di.univaq.it (S. Nocella), fabrizio.rossi@univaq.it (F. Rossi), stefano.smriglio@univaq.it (S. Smriglio).

[1] Telespazio, a Finmeccanica-Alcatel company, is the Italian reference operator for GALILEO system management, navigation and control. The Company has the responsibility for one Mission and Control Constellation Centre, located in Fucino (L'Aquila) and one Signal Performance Validation Centres, located in Rome, as well as the responsibility for the Italian Galileo Test Range; www.telespazio.it

0305-0548/$ - see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.cor.2011.01.016

GALILEO, a joint initiative of the *European Commission* and the *European Space Agency* (ESA), is a global satellite positioning and navigation system (GNSS) specifically designed for civilian purposes, [30]. It is designed to provide worldwide services with outstanding performance in terms of accuracy, integrity, continuity and availability. The whole program cost is about 2.15 billion Euro while the yearly management cost will be about 220 million Euro. The first satellite (Giove-A) of a constellation of 30, was launched in December, 28 2005, while the second satellite (Giove-B) was launched on April, 27 2008. Details on features and current status of GALILEO project can be found in [30].

Major space agencies, such as the *National Aeronautics and Space Administration* [33], the *Indian Space Research Organization* [32], and the *China National Space Administration* [29] developed tools for supporting services scheduling. Descriptions of such tools and relative modeling assumptions can be found, for instance, in [2,19,23,27]. Each tool has specific (mission dependent) constraints that prevent its direct application to similar contexts. This motivated Telespazio and ESA to support the development of an optimization-based decision support tool for service scheduling within the *Innovation Triangle Initiative* framework.[2] This paper illustrates the methodology incorporated within the tool.

The problem of scheduling services on a satellite constellation has been formally introduced in [3] and is referred to as *Satellite Range Scheduling Problem*. This consists of one ground station and a set of services. Each service can only be processed (that is, transmitted) within a given time window, corresponding to the visibility period of the satellite which the service is directed to. Each service has a fixed processing time and, once started, it cannot be interrupted until completion. The ground station can transmit at most one service at a time. The problem consists in finding a schedule that maximizes the number of on-time services, that is, those transmitted within their feasible time window. Even this basic problem turns out to be NP-hard. In fact, Barbulescu et al. [3] show that the *Single Resource Range Scheduling* problem is equivalent to the problem of minimizing the number of *tardy jobs* ($\sum U_j$) on a single-machine with release dates. Finally, the extension in which services are transmitted by two or more ground stations, called *Multi Resource Range Scheduling* problem, is also addressed.

In this paper we study a generalization of the Multi Resource Range Scheduling problem, namely the *Satellite Range Scheduling Problem with Resource Constraints*, (SRSP-RC). This includes several aspects which are quite relevant in practice and, to the best of our knowledge, have never been addressed before in the literature. The first novelty is that, besides ground stations, also satellites have unit transmission capacity, that is, can activate at most one connection at a time. This places the problem into the family of *multiprocessor task scheduling* problems, in which each task (service) requires a set of resources to be processed (ground station and satellite). Furthermore, a set of complex technological constraints are included in the model to faithfully represent the current systems (see Section 2). Finally, a third issue, recommended by ESA and Telespazio engineers, concerns the objective function. In SRSP-RC each service is associated with a profit, and the goal is maximizing the total profit of on-time services. Following the development of [3], this is equivalent to minimizing the weighted sum $\sum w_j U_j$ of tardy jobs.

We propose a time-indexed {0, 1}-linear programming formulation for SRSP-RC able to model the complex structure of the constraints, and show that its linear relaxation provides very good bounds (see Section 7). A major drawback of such a formulation is its size, which often makes even the LP relaxation difficult to solve. Computational advances with time-indexed formulations may often be achieved by Lagrangian relaxation, [1,17,22]. Thus, we devise a strong Lagrangian bound based on a Lagrangian subproblem which extends the one proposed in [1] for the single-machine case, being equivalent to a set of *maximum weighted independent set* problems on *interval graphs*. Unfortunately, the complex structure of the technological constraints prevents the application of standard crossover mechanisms for deriving feasible good schedules within subgradient optimization. Therefore, in Section 6 we propose a heuristic which can be viewed as an adaptation of the *Fix-and-Relax* MIP heuristic developed by Dillenberger et al. [11] to Lagrangian relaxation. The computational experience (Section 7) shows the effectiveness of the approach, yielding near optimal solutions for large scale instances of the GALILEO constellation. The tool incorporating our methodology is currently used by the Telespazio engineers in preliminary simulations of the GALILEO mission.

2. Problem definition and complexity

A *satellite mission* consists of a set $\mathcal{S}=\{1,\ldots,S\}$ of satellites (*space segment*) and a set $\mathcal{M}=\{1,\ldots,M\}$ of ground stations (*ground segment*).

Depending on the mission, satellites may be geosynchronous or not. In the latter case, a ground station m can communicate with a satellite s only when s lies within the transmitting horizon of m. In general, this happens periodically within the planning horizon $\mathcal{T}$. Therefore, for each $(s,m)\in\mathcal{S}\times\mathcal{M}$, the *visibility period* can be expressed as the union of (disjoint) time windows:

$$V(s,m)=\bigcup_{h=1,\ldots,H_{sm}}[t_{sm}^{start(h)},t_{sm}^{end(h)}] \tag{1}$$

Moreover, the system is designed so that each satellite $s\in\mathcal{S}$ is visible from at least one station $m\in\mathcal{M}$ within the time horizon $\mathcal{T}$, that is, $\forall s\in\mathcal{S},\exists m\in\mathcal{M}$ such that $V(s,m)\neq\emptyset$.

Several operations must be performed on spacecrafts, related to either satellite control (e.g., tracking, telemetry reception, telecommand uplink, ranging session, navigation/integrity message uplink), or satellite payload (e.g., earth observation or scientific data download). These operations require ground-to-space communications: we define *service* to be a generic communication between one ground station and one satellite. The information associated with a service is stored in data packets.

The problem under investigation consists of a set of services $\mathcal{J}=\{1,\ldots J\}$ to be scheduled. Each service $j\in\mathcal{J}$ is characterized by the following parameters:

- $s(j)\in\mathcal{S}$, the (unique) satellite requested by j;
- $\mathcal{M}(j)\subseteq\mathcal{M}$, the set of stations able to process j;
- $p_j\in\mathbb{Z}_+$, the (deterministic) processing time, that is, the duration of the communication j;
- $r_j\in\mathcal{T}$, the service release time, i.e., the time in which j becomes available for processing;
- $d_j\in\mathcal{T}$, the service due-date, i.e., the time by which j must be completed;
- $w_j\in\mathbb{Z}_+$, the profit of j.
- $l_j\in\mathbb{Z}_+$, the size of the data packets (Kbytes) associated with j.

[2] ITI supports the identification, validation and development of disruptive space innovations based on new ideas or concepts, giving preference to innovations coming originally from non-space industrial or research sectors. ITI is based on the "Innovation Triangle" concept stating that a rapid and successful introduction of disruptive innovations in Industry requires the collaboration of three different entities: an Inventor as source of innovation; a Developer, casting the new idea into space standards and validating it into a relevant environment; a Customer, that is, a company operating in the space sector that is interested in using the new idea to improve an existing (or create a complete new) space product, see [31].