

Optimal Satellite Payload Selection and Specification

John A. Flory and Jeffrey P. Kharoufeh¹

Department of Industrial Engineering

University of Pittsburgh

1048 Benedum Hall

3700 O'Hara St.

Pittsburgh, PA 15261, USA

Email: jkharouf@pitt.edu

Voice: (412)-624-9832

Final version appears in

Military Operations Research, 15 (3), 2010, pp. 43-57.

Abstract

We consider the problem of optimally selecting and specifying satellite payloads for inclusion on a sequence of satellite launches to maximize the total expected utility of a constellation. The satellite bus is constrained by energy, weight, volume, and cost limitations. We formulate the problem as a variant of the multidimensional knapsack problem and provide exact and heuristic solution methods. Exact solutions are obtained for small and moderately-sized constellations. For larger constellations, the heuristic techniques provide solutions within 4% of the best-known integer solution.

Keywords: Satellite constellation, payload selection, payload specification.

¹Author to whom correspondence should be addressed.

INTRODUCTION

Satellites play a crucial role in a variety of military applications including communications, navigation, imaging, reconnaissance, intelligence gathering, and meteorology. A *satellite constellation* is a collection of satellites working cooperatively to cover a geographical region (or the entire planet) and to accomplish a specific mission. For example, a Global Positioning System (GPS) constellation uses multiple satellites to triangulate an object's precise location on the earth's surface. Typically, satellite constellations are formed and maintained through a schedule of staggered, sequential launches that carry *payloads* into the constellation on a satellite bus. Each payload is designed to perform certain functions over its useful lifetime, and mission planners must decide in advance which capabilities (i.e., which payloads) to include on new satellite launches. These decisions are complicated because (i) different payload types have different levels of importance; (ii) satellite payloads age and deteriorate over time due to the harsh space environment in which they operate; (iii) most satellite payloads are viewed as irreparable systems that must be replaced either when they fail, or when their anticipated useful life has expired; (iv) all payload launch decisions are subject to budget constraints; (v) the selected payloads must be assigned certain engineering specifications to ensure their compatibility with the satellite bus; and (vi) the composition of the existing constellation impacts each launch decision. The primary objective of this research is to determine the optimal set of payloads to include, and specify, on a sequence of satellite launches in order to maximize an overall measure of constellation effectiveness.

When payloads are selected for inclusion on a launch, several specifications are dictated by mission planners (e.g., engineering specifications, functional specifications, resource requirements, reliability specifications, etc.). In this research, we focus on the specification of payload reliability because this measure dictates the anticipated useful lifetime of the payload and its planned replacement time. There are two primary reliability specifications given to payload designers/manufacturers by mission planners: the design life (denoted by L_d) and the mean mission duration (denoted MMD). The design life is the length of time the payload will last with a high probability. For example, if the design life is specified as 10 years with 95% reliability, this means the payload will survive at least 10 years with probability 0.95. Usually, the true lifetime of the unit (call it X) is assumed to follow an exponential distribution (i.e., a constant failure rate model) with survivor (reliability) function $R(t) = P(X > t)$, $t \geq 0$. The MMD is then calculated as the *truncated*

expected lifetime given by (Larson & Wertz, 2004)

$$\text{MMD} = \int_0^{L_d} R(t) dt. \quad (1)$$

However, the method for specifying mean mission duration is not standard in reliability engineering practice. The right-censoring of the survivor function at design life serves as a surrogate for explicitly characterizing the survivor function of items that have increasing failure rates (e.g., consumable propellants and/or cryogenics, mechanical parts, etc.). While the quantity in (1) is, in general, a real number, the payload MMD specification is often stated as an integer value in practice because any further refinement is generally viewed as false precision. Moreover, the fact that most space programs plan to acquire replacement payloads at the MMD reinforces the notion that payloads are considered useful up to this time. For example, if an imaging sensor has been assigned an MMD specification of five years, then the sensor is expected to be operational over this time horizon. At the end of the horizon, the unit must be replaced. Realistically, the sensor may fail prior to the MMD, or it may fail after the MMD. However, the MMD is needed to plan replacements well in advance of their anticipated launch date due to lengthy design and manufacturing lead times. Therefore, the specification of this reliability index is critical in the payload selection and design process.

Payload selection and specification is typically incorporated into the satellite design process wherein the payloads and the satellite bus are designed in tandem. Generally, satellite design is an iterative process in which the satellite bus and payloads are refined, and trade-offs are made between cost, engineering limitations, and performance. Complex “trade spaces” result when different payloads and design parameters are considered. Tools such as the Mission Tradespace Tool (MTT) (Girerd, 2005), used by the Jet Propulsion Laboratory, can assist designers in early identification of mission-feasible spacecraft architectures, and the Advanced Trade Space Visualizer (ATSV) (Yukish et al., 2007), allows a user to obtain Pareto-optimal satellite designs by iteratively passing design parameters to a satellite design model. However, while both the MTT and ATSV can be used to identify designs that meet engineering specifications, they do not consider the uncertainty inherent in satellite and payload performance once they enter service. Richards et al., 2009, present a physics-based model to determine a spacecraft’s beginning-of-life utility and then applies a stochastic model of disturbances to generate lifetime utility trajectories. The model identifies a set of Pareto-optimal designs. Richards’ methodology incorporates the effects of stochastic disturbances on individual satellite performance, but it cannot be used to consider payload selection and design in

the broader context of a satellite constellation. Jilla & Miller, 2004, developed a comprehensive, two-phase methodology to select design parameters for distributed satellite systems. Single and multi-objective optimization models are formulated and solved to optimize constellation performance and life-cycle cost. The decision variables include the number of satellites and various payload design variables. The methodology relies on a Generalized Information Network Analysis (GINA) model of the constellation to evaluate performance by simulating the effects of possible failure states. While this approach is ideally suited for developing the ground-up design of a static satellite constellation, it does not prescribe optimal replacement decisions for an existing constellation. For dynamic constellations comprised of distinct payloads, developing the necessary GINA models as the constellation evolves can be difficult.

In this research, we consider the problem of deciding how best to select and specify satellite payloads for inclusion on a (finite) sequence of satellite launches so that the total expected *utility* of the constellation is maximized. Some clarification of the term *utility* is warranted here. We use the word in a generic sense to mean the benefit (or reward) contributed to the constellation's effectiveness by including satellite payloads (or capabilities). The utility can depend on a number of factors including (but not limited to) mission priorities and the current composition of the constellation. Our objective is not to determine the optimal *timing* of launches but rather the payloads to be included on each launch along with their respective MMD specifications. The utility of each payload is assumed to decrease stochastically over time (once on orbit), and payloads may fail (independently of their utility) prior to reaching their MMD specification. Our model incorporates the notion of *utility dependence* to reflect the reality that the best choice of payloads for the current launch depends explicitly on all prior *and* future launches in the planning horizon. This is a unique feature of our model that complicates the solution procedure considerably. By viewing and solving the problem as a variant of the multidimensional knapsack problem (MKP), we provide near-optimal payload selection and specification decisions for each launch using payload MMD as the integer-valued decision variable. Our problem is analogous to maximizing the overall value of a sequence of MKPs for which the value of each prospective item depends on the past choices as well as the future choices. We provide a binary integer programming (BIP) formulation of the integer knapsack models that simultaneously considers the past and future launch decisions and can be solved exactly using branch-and-cut techniques. In addition, we develop several heuristics to obtain near-optimal solutions for larger constellations with a large number of potential payload types. By way of computational experiments, we assess the effectiveness of the heuristic techniques

by comparing their solutions with those obtained using a commercial solver.

The remainder of the paper is organized as follows. In the next section, we describe the model and introduce the decision variables, objective function, and constraints. Subsequently, we present a variation of the integer multidimensional knapsack problem and an alternative zero-one formulation that accounts for utility dependence. In the solution procedures section, the heuristic solution techniques used to obtain near-optimal solutions are described, while computational experiments are provided in the numerical examples and results section. We conclude the paper with a few closing remarks and discuss directions for future research.

MODEL DESCRIPTION

The satellite constellation is initially (at time 0) populated with S satellites, each of which is equipped with certain payloads. A total of M satellite launches will occur at times n_1, n_2, \dots, n_M such that $0 < n_1 < n_2 < \dots < n_M$, and for each j , $n_j \in \mathbb{Z}_+$, the set of all positive integers. For convenience, let $J = \{1, 2, \dots, M\}$ be the set of satellites (or satellite launches). For each launch $j \in J$, there are K ($K < \infty$) distinct payload types that are candidates for inclusion on the satellite bus. We define the set of payload types as $I = \{1, 2, \dots, K\}$ and assume that no common capabilities exist between two distinct types (i.e., each payload type is responsible for providing a specific capability to the constellation). Furthermore, at most one of each payload type can be included on a given launch. Realistically, payloads in the constellation can operate with degraded capabilities; however, our models assume that the condition of each payload is binary (i.e., either operational or failed) and that random payload lifetimes are mutually independent. The lifetime of payload type $i \in I$ has a non-defective cumulative distribution function (c.d.f.) denoted by F_i .

For our models, the MMD specification serves as both the selection and specification decision variable. Specifically, let x_{ij} denote the MMD specification of payload type $i \in I$ on satellite launch $j \in J$ such that $x_{ij} \in \Theta_i$, for each $j \in J$. That is, the finite set $\Theta_i \subset \mathbb{Z}_+$ is the set of all possible MMD specifications for a type- i payload on each launch. This set always includes the element 0 to allow for the exclusion of payload type i on a given launch. That is, if $x_{ij} = 0$, then payload type i is not selected for inclusion on the j th launch, and if $x_{ij} > 0$, payload type i is selected for inclusion and is assigned MMD specification x_{ij} . Denote by ξ_i the cardinality of Θ_i , i.e. $|\Theta_i| = \xi_i$.

To illustrate these definitions, suppose payload type 2 is a candidate for inclusion on satellite launch 5. The set of possible MMD specifications (in years) for payload type 2 is $\Theta_2 = \{0, 3, 5, 10\}$

so that $\xi_2 = |\Theta_2| = 4$ possible options. For instance, if $x_{25} = 10$, then payload type 2 is included on launch 5 and is expected to be effective in the constellation for 10 years, after which it must be replaced. On the other hand, if $x_{25} = 0$, payload type 2 is not included on launch 5. In this way, the decision variables, x_{ij} , serve as both selection and specification variables.

The objective function can be viewed as an aggregate measure of constellation effectiveness over an appropriate, finite time horizon. Specifically, we consider the maximization of the total expected utility, as defined in the previous section, generated by a sequence of launches. For our models, the utility of a type- i payload at time n is a function of three factors: (i) the relative importance of payload type i ; (ii) the MMD specification of payload type i ; and (iii) the number of operational, type- i payloads in the constellation. We next describe each of these factors and their contribution to the utility function.

First, denote by ψ_i the relative *importance* of payload type i , $i \in I$. Here, ψ_i may assume any positive real number and is not limited to the interval $[0, 1]$. It may be viewed as a measure of the payload's relative functional value to the constellation's mission. Although the importance of a payload type can realistically vary over time, we assume it remains fixed for each type i over the planning horizon. Second, we assume the utility function depends on the MMD specification x_{ij} in the following way. Recall that an increased MMD specification means that the payload is expected to be effective in the constellation for a longer period of time. This may be due to the use of component or subsystem redundancies and/or high-quality construction materials that are less susceptible to degradation. Therefore, we reflect this reality by allowing the form of the utility function to be influenced by the MMD specification. Finally, the utility function depends on the number of functional type- i payloads in the constellation at time n , which is denoted by an integer-valued random variable $Q_n(i)$, $n \geq 0$, $i \in I$. We assume that as the constellation accumulates an increasing number of payloads of the same type, further additions of that type may add only diminishing marginal utility to the constellation. Therefore, on some future launch, it may be optimal to add a different payload type, even if its relative importance is smaller. If we ignore diminishing marginal utility, the problem can be solved by assuming that all the launches are mutually independent. However, this assumption is difficult to justify in practice.

We denote the utility of payload type i on satellite j at time n as $U_n(\psi_i, x_{ij}, Q_n(i))$, $i \in I$, $j \in J$, $n \geq 0$. Note that U_n is random due to its dependence on the (stochastic) number of operational payloads of each type at time n . We do not restrict the form of U_n but only assume that $\mathbb{E}[U_n(\psi_i, x_{ij}, Q_n(i))]$ is a non-increasing function of n and $U_n(\psi_i, 0, Q_n(i)) = 0$ with probability

1 (w.p. 1) for each $n \geq 0$ and $i \in I$. Next, if a type- i payload is included on launch j , we define $\Upsilon_{ij} = \{n_j, n_j + 1, n_j + 2, \dots, n_j + x_{ij} - 1\}$ as the set of years for which the payload contributes to the constellation. For example, if a type- i payload is included on launch j , which occurs at the start of year 2 ($n_j = 2$), and the payload is assigned the MMD value $x_{ij} = 6$, then the payload contributes a benefit to the constellation at the start of each year in the set $\Upsilon_{ij} = \{2, 3, 4, 5, 6, 7\}$. The total expected lifetime utility contributed by a type- i payload on satellite j , denoted by $u(i, j)$, is given by

$$u(i, j) = \begin{cases} \mathbb{E} \left[\sum_{n \in \Upsilon_{ij}} U_n(\psi_i, x_{ij}, Q_n(i)) \right], & \text{if } x_{ij} > 0, \\ 0, & \text{if } x_{ij} = 0. \end{cases} \quad (2)$$

Initially (at time zero), the constellation contains S satellites. Let $J' = \{1, 2, \dots, S\}$ denote the set of existing satellites in the constellation. Because each payload on orbit was assigned a certain MMD specification at the time of its launch, we assume the *remaining MMD* is known for each one. We define the remaining MMD of a type- i payload launched at year $n_{j'}$ as

$$r_{ij'} = n_{j'} + x_{ij'} - n$$

where $x_{ij'}$ denotes the original MMD specification and n is the current year. Let $\bar{u}(i, j')$ be defined as the total expected remaining utility of payload type i on satellite j' , $i \in I$, $j' \in J'$. For example, if the current year is $n = 6$, and a type-3 payload was launched at $n_{j'} = 3$ with $x_{3j'} = 10$, then $r_{3j'} = 7$, and $\bar{u}(3, j')$ is the total expected remaining utility derived from $r_{3j'}$. The remaining utility and MMD specifications of existing payloads will affect the decisions made at launch epochs n_1, n_2, \dots, n_M due to the dependence of U_n on the importance, MMD, and number of operational payloads of the same type. However, the total expected remaining utility of the satellites in the constellation at time zero, given by

$$\bar{u}(i, j') = \mathbb{E} \left[\sum_{n=0}^{r_{ij'}-1} U_n(\psi_i, x_{ij'}, Q_n(i)) \right], \quad (3)$$

does not contain the decision variables x_{ij} , $i \in I$, $j \in J$. The overall objective function over the planning horizon is given by

$$\sum_{i \in I} \sum_{j \in J} u(i, j) + \sum_{i \in I} \sum_{j' \in J'} \bar{u}(i, j'). \quad (4)$$

Finally, we describe the constraints of the optimization model. The energy consumption, weight, volume, and budgetary limitations of each satellite bus must be considered when selecting payloads for each launch. The energy constraint corresponds to the total energy provided by the j th satellite

bus over the planning horizon. This single, aggregate value accounts for (i) the battery lifetime estimate conditioned on the orbit of the satellite and the amount of charging and discharging that takes place due to eclipse and other power conditioning factors; (ii) the expected lifetime of the solar panel which charges the battery; and (iii) other associated power system hardware. We assume that a required minimum threshold of continuous power production (in Watts) is available (i.e., we do not track the available power in each period, but consider total energy as a commodity that is consumed over the planning horizon). The estimate of total available energy is a positive real number, E_j , measured in Watt-years. The maximum allowable total weight of payloads on launch j is denoted by W_j (measured in lbs), and the maximum allowable volume consumed by payloads on launch j is V_j (measured in ft^3). The value V_j is the difference between the volumes of the launch vehicle's payload compartment and the volume of the satellite bus. (Payload geometries are not considered in our formulation.) Finally, for each launch j , there is an associated budget of C_j dollars.

The power consumption (in Watts) of a type- i payload, denoted by A_i , is assumed to be proportional to its utility and independent of the payload's MMD specification. In order to maintain the proportionality of power consumption to utility, we scale the utility function by a factor A'_{ij} for payload i on satellite j given by

$$A'_{ij} = \frac{A_i}{B(\psi_i, x_{ij}, 0)},$$

where $B(\psi_i, x_{ij}, 0)$ is the baseline utility of payload type i on launch j when there are no existing payloads of type i in the constellation. We assume that $0 < B(\psi_i, x_{ij}, 0) < \infty$ w.p. 1. The total expected energy consumed by a type- i payload on satellite j , given a positive MMD specification x_{ij} , is

$$\varepsilon_{ij}(x_{ij}) = \mathbb{E} \left[\sum_{n \in \Upsilon_{ij}} A'_{ij} U_n(\psi_i, x_{ij}, Q_n(i)) \right], \quad i \in I, j \in J. \quad (5)$$

Likewise, the weight, volume, and cost of a payload are modeled as nonnegative, discrete functions of MMD and are nonzero if and only if $x_{ij} > 0$. For $i \in I, j \in J, x_{ij} \in \Theta_i$,

- $w_{ij}(x_{ij})$ is the weight (in lbs) of a type- i payload with MMD specification x_{ij} ;
- $v_{ij}(x_{ij})$ is the volume (in ft^3) occupied by a type- i payload with MMD specification x_{ij} ;
- $c_{ij}(x_{ij})$ is the procurement cost (in dollars) of a type- i payload with MMD specification x_{ij} .

In the next section, we provide an integer multidimensional knapsack formulation for the payload selection and specification problem, as well as an alternative binary integer programming formulation that can be solved to optimality using standard algorithms such as branch-and-cut.

INTEGER KNAPSACK FORMULATIONS

Here, we first formulate the payload selection and specification problem as an integer multidimensional knapsack problem (MKP). From this formulation, we are able to obtain approximate solutions that take into account prior launch decisions but ignore future launch decisions.

Multidimensional Knapsack Model

For a single satellite launch, the payload selection and specification problem can be viewed as a relaxation of the multi-choice, multidimensional knapsack problem (MCMKP). In this variant of the MKP, the set of all items is partitioned, and one item from each partition must be selected (Martello & Toth, 1990). The set of all items, consisting of each unique payload type and MMD specification combination, is

$$\Theta = \bigcup_{i=1}^K \Theta_i. \quad (6)$$

Our problem is a relaxation of the MCMKP because we do not require that one item of each type (or partition) be selected as in the MCMKP. Payloads can be assigned an MMD of zero so, *at most*, one item of each type can be selected. Energy, weight, volume, and cost represent the multiple dimensions of the knapsack. The objective function explicitly includes the contribution to the overall utility measure of payloads initially existing in the constellation. These existing payloads will impact the decisions concerning payloads in all subsequent launches due to the interdependence of utility between payloads of the same type. The multi-satellite, integer knapsack formulation of

the payload selection and specification problem is as follows:

$$\max \sum_{i \in I} \sum_{j \in J} u(i, j) + \sum_{i \in I} \sum_{j' \in J'} \bar{u}(i, j') \quad (7a)$$

$$\text{s.t.} \quad \sum_{i \in I} \varepsilon_{ij}(x_{ij}) \leq E_j, \quad j \in J \quad (7b)$$

$$\sum_{i \in I} w_{ij}(x_{ij}) \leq W_j, \quad j \in J \quad (7c)$$

$$\sum_{i \in I} v_{ij}(x_{ij}) \leq V_j, \quad j \in J \quad (7d)$$

$$\sum_{i \in I} c_{ij}(x_{ij}) \leq C_j, \quad j \in J \quad (7e)$$

$$x_{ij} \in \Theta_i, \quad i \in I, \quad j \in J \quad (7f)$$

where $u(i, j)$ and $\bar{u}(i, j')$ are given by (2) and (3), respectively. Note that the terms $\bar{u}(i, j')$ do not contain the decision variables, x_{ij} . However, the number of type- i payloads in the constellation prior to launch $j = 1$ contributes to $Q_n(i)$, the number of functional type- i payloads at time n , $n \geq 0$.

Were it possible to ignore utility dependence, problem (7) could be solved as a sequence of independent MKPs, each of which ignores the past and future launch decisions. However, viewing the problem in this way is unrealistic and leads to sub-optimal payload selection and specification decisions for the current launch. The dependence structure induces a problem that is analogous to a sequence of MKPs in which the values (and sizes) of items selected for each knapsack depend on the items selected in all prior knapsacks and the future choices. To complicate matters further, formulation (7) cannot be solved directly using standard integer programming (IP) solution techniques because the utility functions, $u(i, j)$, and the energy functions, $\varepsilon_{ij}(x_{ij})$, are implicit functions of the decision variables, $x_{i\ell}$, where $\ell \in J$ and $\ell < j$. Therefore, formulation (7) violates the additivity axiom of integer programming which requires the objective function and constraints to be additively separable in the decision variables. Consequently, standard IP solution techniques (e.g., the branch-and-cut method) that solve the linear programming relaxation of the problem cannot be applied. However, approximate solutions to problem (7) that incorporate dependence on *prior* launches can be obtained using heuristic procedures (as we demonstrate in the next two sections). In the next subsection, we propose an alternative binary integer programming (BIP) formulation that circumvents these complications and provides the optimal payload selection and MMD specification decisions over all M launches simultaneously.

BIP Formulation

Here, we present a BIP formulation in which the total expected utility and energy consumption are not implicit functions of the decision variables. Instead, they can be viewed as constants that are additively separable in the decision variables. To this end, we consider the MMD specifications of each payload type on all M satellite launches, and in this formulation, payload specification is viewed as the distinct ways in which the MMDs of each payload type can be allocated among the M satellite launches. The number of permutations of MMD specifications for payload type i among the M satellites is ξ_i^M , $i \in I$. Define the set $Z_i = \{1, 2, \dots, \xi_i^M\}$. The MMD specification of payload type i on satellite j in the k th permutation of Z_i is denoted by $x_{ij}^{(k)}$, $i \in I$, $j \in J$, $k \in Z_i$. For example, consider a type-1 payload whose MMD specification is to be determined for $M = 4$ satellite launches, and suppose it can be assigned one of four distinct MMD specifications (i.e., $\xi_1 = 4$). In this case, the set Z_1 contains $4^4 = 256$ possible MMD specification permutations. That is, there are 256 different ways to assign MMD to payload type 1 over the four launches.

We map each permutation in Z_i to a binary decision variable, y_{ik} , such that

$$y_{ik} = \begin{cases} 1, & \text{if } [x_{i1}^{(k)}, x_{i2}^{(k)}, \dots, x_{iM}^{(k)}] \text{ is chosen,} \\ 0, & \text{otherwise} \end{cases}$$

for $i \in I$ and $k \in Z_i$. The total expected lifetime utility associated with a permutation of MMD for payload type i over all M satellite launches is the sum of the M total expected lifetime utilities of each type- i payload. Let $u^{(k)}(i, j)$ be the total expected lifetime utility of payload type i on satellite j in the k th permutation. By summing over all M launches, we obtain the total utility of payload type i denoted by

$$u_{ik} = \sum_{j \in J} u^{(k)}(i, j), \quad i \in I, k \in Z_i \quad (8)$$

where $u^{(k)}(i, j)$ is computed by equation (2) for each $k \in Z_i$. Moreover, for the k th permutation of payload type i , let the total consumption of energy, weight, volume, and cost resources on satellite j be denoted by constants $\varepsilon_{ij}^{(k)}$, $w_{ij}^{(k)}$, $v_{ij}^{(k)}$, and $c_{ij}^{(k)}$, respectively, where $\varepsilon_{ij}^{(k)} \equiv \varepsilon_{ij}(x_{ij}^{(k)})$, $w_{ij}^{(k)} \equiv w_{ij}(x_{ij}^{(k)})$, $v_{ij}^{(k)} \equiv v_{ij}(x_{ij}^{(k)})$, and $c_{ij}^{(k)} \equiv c_{ij}(x_{ij}^{(k)})$. As before, E_j, W_j, V_j , and C_j are the energy, weight, volume, and cost resources associated with satellite launch $j \in J$.

The revised BIP formulation is as follows:

$$\max \sum_{i \in I} \sum_{k \in Z_i} u_{ik} y_{ik} \quad (9a)$$

$$\text{s.t.} \quad \sum_{i \in I} \sum_{k \in Z_i} \varepsilon_{ij}^{(k)} y_{ik} \leq E_j, \quad j \in J \quad (9b)$$

$$\sum_{i \in I} \sum_{k \in Z_i} w_{ij}^{(k)} y_{ik} \leq W_j, \quad j \in J \quad (9c)$$

$$\sum_{i \in I} \sum_{k \in Z_i} v_{ij}^{(k)} y_{ik} \leq V_j, \quad j \in J \quad (9d)$$

$$\sum_{i \in I} \sum_{k \in Z_i} c_{ij}^{(k)} y_{ik} \leq C_j, \quad j \in J \quad (9e)$$

$$\sum_{k \in Z_i} y_{ik} = 1, \quad i \in I \quad (9f)$$

$$y_{ik} \in \{0, 1\}, \quad i \in I, \quad k \in Z_i. \quad (9g)$$

For an M -launch, K -payload type specification problem in which each payload has ξ candidate MMD specifications, formulation (9) has $K\xi^M$ binary variables and $K+4M$ constraints. In general, it is difficult to solve integer programs for which the number of variables greatly exceeds the number of constraints. However, the K equality constraints of (9f) are special-ordered sets of type one (SOS1) constraints that allow only one of the y_{ik} variables to be nonzero in each set. The presence of these constraints ensures that the optimal solution(s) of (9) can (eventually) be obtained. In the next section, we discuss solution methods for problems (7) and (9). Specifically, we obtain optimal (or near-optimal) solutions of (9) using a commercial solver, and we propose four heuristic techniques to obtain approximate optimal solutions of (7).

SOLUTION PROCEDURES

For very small problems, the BIP formulation (9) can be solved by exhaustive enumeration (and checking for energy, cost, weight, and volume feasibility). However, even a modestly-sized problem has a large solution space. For example, an instance with $M = 5$ launches, $K = 4$ payload types, and $\xi = 4$ candidate MMD specifications has $2^{40} \approx 10^{12}$ candidate solutions. To obtain optimal (and near-optimal) solutions to problem (9), we used Dash Optimization's commercial solver Xpress-MP. Like most commercial solvers, Xpress-MP employs the branch-and-cut method (Mitchell, 2002) which uses both branch-and-bound and cutting-plane algorithms in concert. Typically, these algorithms are augmented with additional preprocessing and search heuristics. The branch-and-cut method is guaranteed to converge to the optimal solution (if it exists) of a BIP;

however, computational constraints (e.g., limited memory resources) may cause the method to terminate prematurely with a “best-known” solution.

In order to assess the impact of utility dependence, and to improve computation time, we also propose four heuristics to obtain approximate optimal solutions to problem (7). Greedy solutions to this problem are obtained by solving a sequence of M multidimensional knapsack problems, each corresponding to one of the M launches, that ignores the future launch decisions. That is, we consider the effect of all *prior* launches on the current launch, but not the effect of future launches.

Norm-Based Heuristics

The motivation for developing norm-based heuristics is the classical profit-to-weight ratio heuristic (Dantzig, 1957) for the one-dimensional knapsack problem. However, for individual satellite buses, the one-dimensional profit-to-weight ratio heuristic cannot be directly applied to the payload selection problem because the satellite bus is a four-dimensional knapsack. Therefore, we propose a single, scalar quantity that serves as an aggregate measure of resource consumption. The scalar is constructed in the spirit of the Toyoda heuristic for MKPs (Toyoda, 1975) which views an item’s worth in terms of all the resources it consumes. By scaling each payload type’s utility by this value, we obtain an analogous profit-to-weight ratio.

Consider a single launch j , and define the set of all unique payload type/MMD pairs,

$$\Pi_j \equiv \{(i, x_{ij}) : i \in I, x_{ij} \in \Theta_i\}, \quad j \in J.$$

For simplicity, we assume each payload type has the same number of candidate MMD specifications (i.e., $|\Theta_i| = \xi$ for each $i \in I$). For each launch, there are $d = K\xi$ such payload type/MMD pairs, so we define an index set $\Delta = \{1, 2, \dots, d\}$. Denote by \hat{u}_{mj} the overall utility generated by combination m on launch j , and define $\hat{\varepsilon}_{mj}$, \hat{w}_{mj} , \hat{v}_{mj} , and \hat{c}_{mj} as the energy, weight, volume, and cost resources used by combination m on launch j , respectively. For combination m on launch j , define the vector

$$\omega'_{mj} = [\hat{\varepsilon}_{mj}/E_j, \hat{w}_{mj}/W_j, \hat{v}_{mj}/V_j, \hat{c}_{mj}/C_j].$$

That is, ω'_{mj} is the vector of ratios of the payload’s resource requirements to the satellite bus’s resource capacities, $m \in \Delta$, $j \in J$. For a positive integer p , the p -norm of a vector $y = [y_1, y_2, \dots, y_n]$, is defined as (Meyer, 2000)

$$\|y\|_p = (y_1^p + y_2^p + \dots + y_n^p)^{1/p}. \quad (10)$$

Our aggregate measure of payload resource consumption is the p -norm of ω'_{mj} given by

$$\omega_{mj} \equiv \|\omega'_{mj}\|_p, \quad m \in \Delta, j \in J. \quad (11)$$

For launch j , we compute the quantities \hat{u}_{mj} and ω_{mj} , and their ratio \hat{u}_{mj}/ω_{mj} . For $m, n \in \Delta$, we say combination m has priority over combination n if

$$\frac{\hat{u}_{mj}}{\omega_{mj}} \geq \frac{\hat{u}_{nj}}{\omega_{nj}}.$$

Payloads are added to the satellite bus greedily (i.e., in order of decreasing ratio), subject to the resource constraints and the constraint that, at most, one of each payload type can be included. Ties are broken arbitrarily. The p -norm heuristic was empirically tested on randomly-generated problem instances for various integer values of p . The results of these experiments indicated that the 5-norm heuristic performed well consistently; therefore, we set $p = 5$ in all of the problem instances of the next section. By using a non-Euclidean norm, the value of ω_{mj} is most influenced by those resources whose capacities are most consumed. The next heuristic provides a way to consider relative resource consumption.

Weighted-Norm (WN) Heuristic

While ω_{mj} represents a measure of resource requirements for a particular payload type/MMD combination, it does not take into account the relative scarcity of each resource. For example, if a satellite bus is severely limited in its energy capacity, it is intuitive that payloads having a larger relative energy requirement should have a greater effect on the norm's value. Let $S_\varepsilon(j)$, $S_w(j)$, $S_v(j)$, and $S_c(j)$ represent the relative scarcity of energy, weight, volume, and cost resources on satellite j , respectively, so that for each $j \in J$,

$$\begin{aligned} S_\varepsilon(j) &= E_j^{-1} \sum_{m \in \Delta} \hat{\varepsilon}_{mj}, \\ S_w(j) &= W_j^{-1} \sum_{m \in \Delta} \hat{w}_{mj}, \\ S_v(j) &= V_j^{-1} \sum_{m \in \Delta} \hat{v}_{mj}, \\ S_c(j) &= C_j^{-1} \sum_{m \in \Delta} \hat{c}_{mj}. \end{aligned}$$

For each payload combination $m \in \Delta$ and each launch $j \in J$, define a weighted 2-norm, μ_{mj} , by

$$\mu_{mj} = [S_\varepsilon(j)(\hat{\varepsilon}_{mj}/E_j)^2 + S_w(j)(\hat{w}_{mj}/W_j)^2 + S_v(j)(\hat{v}_{mj}/V_j)^2 + S_c(j)(\hat{c}_{mj}/C_j)^2]^{1/2}. \quad (12)$$

The scalar μ_{mj} is the counterpart to ω_{mj} in the 5-norm heuristic. This method is called the *weighted norm* heuristic. For launch j , we compute the quantities \hat{u}_{mj} and μ_{mj} , and their ratio \hat{u}_{mj}/μ_{mj} . For $m, n \in \Delta$, we say combination m has priority over combination n if

$$\frac{\hat{u}_{mj}}{\mu_{mj}} \geq \frac{\hat{u}_{nj}}{\mu_{nj}}.$$

Payloads are added to the satellite bus greedily (i.e., in order of decreasing ratio), subject to the resource constraints and the constraint that, at most, one of each payload type can be included. Ties are broken arbitrarily.

Greedy Heuristic

The performance of the 5-norm and weighted-norm heuristics is compared to that of a purely greedy heuristic that bases selection decisions solely on the total utility of payload type/MMD combinations. While all the heuristic methods presented here are “greedy” (i.e., they maximize the total expected utility of each launch successively without regard for future launches), the greedy heuristic is so named because it selects payload type/MMD combinations based only on the magnitude of their total expected utility, irrespective of resource consumption. The greedy heuristic first computes the utility \hat{u}_{mj} for each $m \in \Delta$ and $j \in J$. Then, for $m, n \in \Delta$, combination m has priority over n if

$$\hat{u}_{mj} \geq \hat{u}_{nj}.$$

As in the norm-based heuristics, the greedy heuristic attempts to include payloads and assign MMD in order of decreasing utility value, subject to the feasibility constraints and the requirement that at most one of each payload type can be included on a single launch. Ties are broken arbitrarily.

Simulated Annealing Heuristic

The Simulated Annealing (SA) heuristic, like the norm-based heuristics, seeks to maximize the utility of each satellite launch sequentially, so each launch is viewed as an individual MKP that takes into account the decisions made on all prior launches, irrespective of future launches. For each $m \in \Delta$, define the binary variable x_m by

$$x_m = \begin{cases} 1, & \text{if combination } m \text{ is selected,} \\ 0, & \text{otherwise.} \end{cases}$$

A solution $x = [x_1, x_2, \dots, x_d] \in \{0, 1\}^d$ is feasible if (i) for each $i \in I$, at most one type- i payload is included, and (ii) the energy, weight, volume, and cost constraints of launch j are not violated. For a given solution x , a neighboring solution, x' , is any solution that differs from x by the inclusion or exclusion of one payload type/MMD combination. Therefore, a *move* is defined by adding or removing a payload, subject to feasibility. Recall that, in the SA heuristic, T_0 is the initial temperature, and T_f is the terminal temperature, where $T_f < T_0$. We use a geometric cooling schedule with cooling ratio $r \in (0, 1)$. At each temperature level, N solutions are explored.

As noted in (Michalewicz & Fogel, 2000), selection of the parameters T_0 , T_f , r , and N is integral to the performance of the SA heuristic. SA procedures are known to perform well for high initial temperatures and sufficiently slow cooling schedules (Park & Kim, 1998). Therefore, an initial temperature, $T_0 = 500$, was selected so that the initial probability of accepting a neighboring solution was close to 1.0. A cooling ratio, $r = 0.05$, was used with a final temperature of $T_f = 0.05$, and $N = 5$ neighboring solutions were explored at each temperature. For our problems, SA exhibits strong dependence on the initial solution; therefore, for each problem instance, 30 independent replications were performed with distinct initial solutions.

In the next section, we assess the quality of solutions obtained by the norm-based, weighted-norm, greedy, and SA heuristics to problem (7). These solutions are compared to optimal (or near-optimal) solutions generated by solving (9) which accounts for prior and future launches simultaneously.

NUMERICAL EXAMPLES AND RESULTS

In this section, we present numerical experiments to compare the performance of the exact and heuristic solution methods using realistic payload data. Performance is evaluated by two attributes: solution quality and computation time (in seconds). Although Xpress-MP Solver uses a branch-and-cut technique that is guaranteed to (eventually) converge to the optimal solution of a BIP, we report the best solution obtained in cases where computational limits cause premature termination of the branch-and-cut scheme. Each of the four heuristics were coded in the MATLAB computing environment, and all experiments were performed on a personal computer equipped with a 2.66 GHz Intel Xeon processor with 2.00 GB of RAM. To solve instances of (7) and (9), a functional form for U_n and realistic problem data describing payload and satellite bus attributes are required. These two issues are discussed next.

Payload Utility Function

It is important to note that the integer knapsack formulations make no assumptions about the form of the payload utility function U_n . Here, we define a utility function that can be used to characterize payload behavior over time. Let $\mathbf{1}_n(i, j)$ be an indicator function denoting the functional status at time n of a type- i payload included on launch j . Specifically, $\mathbf{1}_n(i, j) = 1$ if the payload is operational at time n , and $\mathbf{1}_n(i, j) = 0$ otherwise. For each $n > n_j$ and $x_{ij} > 0$, the payload utility function, U_n , conditioned on $\mathbf{1}_n(i, j)$, is given by

$$U_n(\psi_i, x_{ij}, Q_n(i) | \mathbf{1}_n(i, j) = k) = \begin{cases} \frac{\psi_i e^{-\beta(n-n_j)/x_{ij}}}{\max\{q_n^\gamma(i), 1\}} + N(0, \sigma_{n-n_j}^2), & k = 1, \\ 0, & k = 0, \end{cases} \quad (13)$$

where $q_n(i) \equiv \mathbb{E}[Q_n(i)]$, the parameter β ($\beta > 0$) adjusts the rate of utility decrease, γ ($\gamma \geq 0$) is a utility dependence parameter (described below), and $N(0, \sigma_{n-n_j}^2)$ denotes a normally distributed noise term with mean zero and variance $\sigma_{n-n_j}^2$ for $n > n_j$.

The term $\max\{q_n^\gamma(i), 1\}$ is used to model diminishing marginal utility so that, as the expected number of operational type- i payloads in the constellation increases, the contribution of an additional type- i payload to the overall utility diminishes. The utility dependence constant γ is chosen such that $\gamma = 0$ implies independence, and $\gamma > 0$ implies some measure of utility dependence between like payloads. We take the maximum of $q_n^\gamma(i)$ and 1 so that (13) is well defined if $q_n(i)$ is close to zero. Taking the expectation of both sides of (13), we obtain

$$\mathbb{E}[U_n(\psi_i, x_{ij}, Q_n(i) | \mathbf{1}_n(i, j) = k)] = \begin{cases} \frac{\psi_i e^{-\beta(n-n_j)/x_{ij}}}{\max\{q_n^\gamma(i), 1\}}, & k = 1, \\ 0, & k = 0. \end{cases} \quad (14)$$

This functional form implies that the expected utility of a functioning payload (of type i) decreases exponentially over time.

For illustrative purposes, we assume that each payload type has a constant failure rate (i.e., an exponential lifetime distribution), and that all payload lifetimes are mutually independent. (Note that the constant failure rate assumption can be relaxed since only the lifetime c.d.f., F_i , is needed.) The probability that a payload of type i is operational at time n , given that it was launched at time n_j , is

$$\mathbb{P}(\mathbf{1}_n(i, j) = 1) = 1 - F_i(n - n_j) = e^{-\alpha(n-n_j)/x_{ij}}, \quad n > n_j.$$

That is, the failure rate of a type- i payload included on launch j and assigned MMD specification x_{ij} is α/x_{ij} with $\alpha > 0$ and $x_{ij} > 0$. Next, we compute the unconditional expectation of U_n for use

in equation (2). For $n > n_j$, we obtain

$$\begin{aligned}
\mathbb{E}[U_n(\psi_i, x_{ij}, Q_n(i))] &= \sum_{k=0}^1 \mathbb{E}[U_n(\psi_i, x_{ij}, Q_n(i) | \mathbf{1}_n(i, j) = k)] \mathbb{P}(\mathbf{1}_n(i, j) = k) \\
&= \mathbb{E}[U_n(\psi_i, x_{ij}, Q_n(i) | \mathbf{1}_n(i, j) = 1)] \mathbb{P}(\mathbf{1}_n(i, j) = 1) \\
&= \frac{\psi_i}{\max\{q_n^\gamma(i), 1\}} \exp\left[\frac{-(\alpha + \beta)(n - n_j)}{x_{ij}}\right].
\end{aligned}$$

Therefore, for $x_{ij} > 0$, $u(i, j)$ is given by

$$\begin{aligned}
u(i, j) = \mathbb{E}\left[\sum_{n \in \Upsilon_{ij}} U_n(\psi_i, x_{ij}, Q_n(i))\right] &= \sum_{n \in \Upsilon_{ij}} \mathbb{E}[U_n(\psi_i, x_{ij}, Q_n(i))] \\
&= \sum_{n \in \Upsilon_{ij}} \frac{\psi_i}{\max\{q_n^\gamma(i), 1\}} e^{-(\alpha + \beta)(n - n_j)/x_{ij}}.
\end{aligned}$$

For each numerical illustration, we used the following parameter values: $\alpha = |\ln 0.9|$, $\beta = |\ln 0.5|$, and $\gamma = 0.5$. These values imply that payloads will survive to their MMD with 90% probability, and upon reaching their MMD, they will contribute 50% of their original utility. For example, a type- i payload included on satellite launch j will survive to period $n = n_j + x_{ij}$ with probability 0.90, at which time its expected utility will be diminished by 50%.

Problem Data

Realistic ranges for payload and satellite bus attributes were obtained from a systems engineer with over 20 years experience in satellite planning. Generally, a payload's cost is proportional to its weight; therefore, cost and weight values are correlated. We generated data for eight distinct payload types, where each type can be assigned an MMD specification of 0, 3, 6, or 10 years, i.e., $\Theta_i = \{0, 3, 6, 10\}$ for $i = 1, 2, \dots, 8$. The same set of MMD choices was used for each payload type to facilitate comparison of results, but in general, $\Theta_i \neq \Theta_j$ for $j \neq i$. Payloads that might include redundant systems, or that are constructed using high-quality materials, will generally be more costly and have higher weight and volume requirements. Therefore, $c_{ij}(x_{ij})$, $w_{ij}(x_{ij})$, and $v_{ij}(x_{ij})$ are proportional to the MMD specification, x_{ij} . The importance of a type- i payload (ψ_i) is roughly proportional to its aggregate resource requirements. This assumption prevents a trivial solution in which a payload with high importance and low resource requirements is always selected. The data used for the first set of numerical illustrations are summarized in Table 1.

Table 1: Summary of payload parameter values.

Type (i)	x_{ij} (yrs)	ψ_i	A_i (Watts)	$c_{ij}(x_{ij})$ (\$100K)	$w_{ij}(x_{ij})$ (lbs)	$v_{ij}(x_{ij})$ (ft ³)
1	3	10.0	500	425	450	15.0
	6	10.0	500	460	475	17.0
	10	10.0	500	500	500	20.0
2	3	8.5	475	375	400	16.0
	6	8.5	475	405	415	18.0
	10	8.5	475	430	420	19.5
3	3	7.5	425	410	430	10.0
	6	7.5	425	460	480	13.0
	10	7.5	425	480	495	14.0
4	3	7.0	260	300	230	10.0
	6	7.0	260	350	280	13.0
	10	7.0	260	370	300	14.0
5	3	6.0	225	370	380	13.0
	6	6.0	225	400	390	15.5
	10	6.0	225	410	395	17.5
6	3	5.5	300	280	240	8.0
	6	5.5	300	320	290	9.0
	10	5.5	300	380	310	12.0
7	3	5.0	275	150	280	7.0
	6	5.0	275	190	350	9.5
	10	5.0	275	240	410	14.0
8	3	3.0	175	270	225	4.0
	6	3.0	175	310	360	5.5
	10	3.0	175	335	300	8.0

Ranges were also obtained for the resource capacities of the satellite buses. It is assumed throughout that all satellite buses have identical resource capacities, so each satellite's loading can be compared more easily. The satellite bus data are provided in Table 2.

A variety of problem instances were generated, both simple and complex. We define problem

Table 2: Assumed satellite bus limitations.

Energy Capacity (W-yrs)	Weight (lbs)	Volume (ft ³)	Budget (\$100K)
10,500	2,500	100	2,500

complexity by three primary variables: the number of satellite launches, the number of distinct payload types, and the number of candidate MMD specifications. Each payload type here is assumed to have four candidate MMD specifications to ensure realism. The number of satellite launches and payload types, however, are varied. Problem size is determined by the number of binary (0-1) variables in the BIP formulation (9). Small problems have, at most, 100 binary variables; medium problems have 101 – 10,000 binary variables; and large problems have more than 10,000 binary variables. Table 3 summarizes the small, medium, and large instances considered in the first experiment.

Table 3: Payload selection problem instances.

Problem Size	Attribute	Value
Small	Number of launches	2
	Number of payload types	5
	Number of binary variables	80
Medium	Number of launches	4
	Number of payload types	6
	Number of binary variables	1,536
Large	Number of launches	7
	Number of payload types	8
	Number of binary variables	131,072

Satellite launches are assumed to occur at the start of years 2, 5, 7, 9, 10, 11, 13, and 15 so that $n_1 = 2, n_2 = 5, n_3 = 7, \dots, n_8 = 15$. These relatively short inter-launch intervals are designed to induce utility dependence. All problem instances are initiated by selecting payloads for launch at time $n_1 = 2$.

Results and Summary

This section summarizes the performance results of the problem instances described in Table 3. Each heuristic procedure is applied to problem (7) while Xpress-MP Solver is applied to problem (9). Table 4 summarizes the objective function values, relative performance of the techniques, and computation time for the small problem instance.

Table 4: Maximum total expected utility (small instance).

Method	Total Expected Utility	% Diff. Optimal	Computation Time (sec)
Xpress-MP Solver	437.6	0.0	0.70
Simulated Annealing	435.1	0.6	20.76
Greedy Heuristic	405.2	7.4	0.07
5-Norm	430.1	1.7	0.14
Weighted Norm	430.1	1.7	0.11

Xpress-MP Solver required only 0.70 sec to obtain the optimal solution. The performance of the two norm-based heuristics was comparable, and they were closer to optimal than the greedy heuristic. Both norm-based heuristics and SA obtained solutions within 2% of optimality. While SA obtained a slightly better solution than the norm-based heuristics, it required significantly more time. The optimal, or near-optimal, solutions obtained by each procedure are summarized in Table 5.

Table 5: Best attainable solutions (small instance).

Method	Xpress-MP		SA		GH		5-Norm		WN	
	1	2	1	2	1	2	1	2	1	2
Payload 1 MMD	10	10	10	10	10	10	10	10	10	10
Payload 2 MMD	3	10	3	10	10	10	6	10	10	3
Payload 3 MMD	3	3	6	3	0	10	0	3	0	10
Payload 4 MMD	10	10	6	10	10	0	10	10	10	10
Payload 5 MMD	10	10	10	10	3	6	10	10	3	10

We note that all of the procedures include payload type 1 on both launches and assign the highest possible MMD specification to this payload. The result is intuitive because the utility function is increasing in both ψ_i and the MMD specification, and payload type 1 has the highest importance

level. The solution obtained by the SA heuristic was very similar to the optimal solution; however, it provided different optimal specifications for payload types 3 and 4. As expected, the greedy heuristic roughly attempts to assign the highest possible MMD value to payloads in order of their importance.

Table 6 summarizes the results of the medium problem instance. For this problem, Xpress-MP required a computation time nearly two orders of magnitude greater than that needed for the small instance to obtain the optimal solution.

Table 6: Maximum total expected utility (medium instance).

Method	Total Expected Utility	% Diff. Optimal	Computation Time (sec)
Xpress-MP Solver	874.5	0.0	85.00
Simulated Annealing	850.1	2.8	47.90
Greedy Heuristic	815.7	6.7	0.10
5-Norm	835.4	4.5	0.11
Weighted Norm	822.7	5.9	0.15

SA obtained a solution within 3% of optimality in approximately half the time. Of the remaining heuristics, only the 5-norm obtained a solution within 5% of optimality.

For the large problem instance in Table 7, it was not possible to solve the BIP to optimality due to RAM limitations. Instead, Xpress-MP reported the duality gap of the BIP, providing upper and lower bounds on its optimal objective value. The lower bound corresponds to the best integer solution obtained, and the upper bound is determined by the cutting planes generated by the branch-and-cut algorithm. Therefore, the heuristic solutions are compared to the best integer solution obtained, as well as the least upper bound of the optimal objective function value. The 5-norm heuristic obtained a solution that was closest to the best solution obtained, followed by the weighted-norm heuristic.

Table 7: Maximum total expected utility (large instance).

Method	Total Expected Utility	% Diff. Best Solution	% Diff. Upper Bound	Computation Time (sec)
Xpress-MP Solver	1469.7	0.0	2.8	7012.70
Simulated Annealing	1415.6	3.7	6.3	151.88
Greedy Heuristic	1420.4	3.4	6.0	0.14
5-Norm	1441.3	1.9	4.6	0.16
Weighted Norm	1426.5	2.9	5.6	0.21

The three problem instances illustrate that, as the number of launches and candidate MMD specifications increase, the computational effort required to obtain exact solutions to problem (9) increases substantially, and for large problem instances, exact solutions may not be attainable. Of the heuristic techniques used to solve problem (7), both the 5-norm and SA heuristics obtained solutions that were within 5% of optimality for the small and medium instances, and within 4% of the best-known integer solution for the large instance. The two norm-based and SA heuristics performed better than the greedy heuristic in every problem instance. The 5-norm heuristic exhibited the best overall performance considering both solution quality and computation time for the three problem instances.

Randomly-Generated Problem Instances

In this section, randomly-generated instances of problem (7) are solved using each heuristic, and their solutions are compared to optimal solutions of problem (9) obtained by the Xpress-MP Solver. A single-launch, eight-payload type problem is considered to compare heuristic performance on a problem analogous to a single instance of an MKP. The parameter values for payload and satellite bus attributes were drawn from uniformly distributed populations, and the payload importance measures are identical to those given in Table 1. Payloads of the same type are assigned the same energy requirement; however, there is no induced correlation between payload MMD and randomly-generated resource consumption.

The power requirement (in Watts) of each payload is drawn from a $U(100, 500)$ population, the cost (in \$100K) of each payload type is drawn from a $U(100, 500)$ population, the payload weight (in lbs) comes from a $U(100, 500)$ population, and the payload volumes (in ft^3) are uniformly

distributed on [3, 20]. The satellite bus has a maximum energy capacity (in Watt-yr) that is drawn from a $U(7500, 12500)$ population, the maximum allowable weight (in lbs) is $U(2000, 3000)$, the maximum allowable volume (in ft^3) is $U(75, 125)$, and the maximum available budget (in \$100K) is $U(2000, 3000)$.

We generated 100 independent instances of problem (9) and solved each to optimality. The mean, standard deviation, maximum, and median of the difference between the optimal objective function value and the value obtained by each heuristic were computed. The mean and median differences illustrate the overall accuracy of each heuristic with more and less sensitivity, respectively, to outlying differences. The standard deviation and maximum differences are used to evaluate how consistently each heuristic obtains near-optimal solutions and to determine worst-case performance. The heuristics' performance on the randomly-generated problem instances are summarized in Table 8. The mean differences from optimality of the heuristics' solutions were similar. SA has the lowest mean difference followed by the 5-norm with a slightly higher mean difference. The lower standard deviation and significantly lower maximum difference from optimality of SA indicate that it provides near-optimal solutions more consistently than do the other heuristics; however, it requires substantially more computation time. The weighted-norm heuristic performed roughly equivalently to the greedy heuristic.

Table 8: Heuristic solution quality (100 problem instances).

Method	% Mean Diff.	% Std. Dev.	% Max Diff.	% Median Diff.	Time (sec)
Simulated Annealing	2.4931	2.0204	7.7330	2.1342	20.87
Greedy Heuristic	5.0545	5.1807	24.9068	4.3784	0.02
5-Norm	3.0781	3.5238	15.6261	1.8375	0.02
Weighted Norm	4.9201	5.4404	26.9841	3.3097	0.02

CONCLUSIONS

In this research, we have formulated and solved optimization problems to determine the optimal, or near-optimal, set of payloads to include and specify on a sequence of satellite launches in order to maximize an overall measure of constellation effectiveness. These models can help mission planners by providing a structured decision-making framework that allows for flexibility in modeling payload lifetimes and the utility function. Moreover, this approach is widely applicable to different types of

satellite constellations because it uses characteristics that are common to all payloads and satellite buses.

While the models and solution techniques are relatively easy to implement, there are a few important shortcomings of the models that are noteworthy. First, we assumed that the sequence of launch times was known in advance. However, in many applications, the timing of satellite launches depends on a number of factors including budgetary limits, time-varying mission priorities, and payload or satellite bus manufacturing lead times. Determining the optimal timing of satellite launches is an interesting problem in its own right. Second, although our models considered the volume requirements of individual payloads, their specific geometries were not considered. The solutions obtained by our models must still be evaluated by design engineers to determine their feasibility with respect to geometric, thermal, or other engineering constraints related to the satellite bus. In the future, it will be instructive to work closely with satellite engineers to establish an expanded set of design constraints, and to develop utility functions that accurately reflect current operations and mission priorities.

Acknowledgements

The authors thank three anonymous referees and Dr. Dick Deckro for useful comments and suggestions that have improved the relevance and presentation of this work. We also acknowledge, with gratitude, Mr. Justin Comstock who introduced us to the payload selection and specification problem.

References

- [1] Dantzig, G. 1957. Discrete variable extremum problems. *Operations Research*, Vol 5, 266-277.
- [2] Girerd, A. 2005. A rapid, flexible approach to tradespace definition and exploration. In *2005 IEEE Aerospace Conference*, March 5-12, 2005, 4212-4219.
- [3] Jilla, C.D. and Miller, D.W. 2004. Multi-objective, multidisciplinary design optimization methodology for distributed satellite systems. *Journal of Spacecraft and Rockets*, Vol 41, No 1, 39-50.
- [4] Keeney, R. and Raiffa, H. 1993. *Decisions with Multiple Objectives*. Cambridge University Press, 141-148.

- [5] Larson, W. and Wertz, J. 2004. *Space Mission Analysis and Design*. Microcosm Press, 767.
- [6] Martello, S. and Toth, P. 1990. *Knapsack Problems: Algorithms and Computer Implementations*. John Wiley & Sons, Inc., 317.
- [7] Meyer, C. 2000. *Matrix Analysis and Applied Linear Algebra*, SIAM, 274.
- [8] Michalewicz, Z. and Fogel, D. 2000. *How to Solve It: Modern Heuristics*. Springer-Verlag, 117-125.
- [9] Mitchell, J. 2002. Branch-and-cut algorithms for combinatorial optimization problems. *Handbook of Applied Optimization*, Oxford University Press, 65-68.
- [10] Park, M. and Kim, Y. 1998. A systematic procedure for setting parameters in simulated annealing algorithms. *Computers and Operations Research*, Vol 25, 207-217.
- [11] Richards, M.G., Ross, A., Shah, N., and Hastings, D. 2009. Metrics for evaluating survivability in dynamic multi-attribute tradespace exploration. *Journal of Spacecraft and Rockets*, Vol 46, No 5, 1049-1064.
- [12] Toyoda, Y. 1975. A simplified algorithm for obtaining approximate solutions to zero-one programming problems. *Management Science*, Vol 21, No 12, 1417-1427.
- [13] Yukish, M., Stump, G.M. and Lego, S. 2007. Visual steering and trade space exploration. In *2007 IEEE Aerospace Conference*, March 3-10, 2007, 1-9.