

# Revenue Management for Communication Satellite Operators – Opportunities and Challenges

Markus Guerster\*, Joël Grotz#, Peter Belobaba\*, Edward Crawley\*, Bruce Cameron\*

\* Massachusetts Institute of Technology  
77 Massachusetts Ave 33-409  
Cambridge, MA 02139  
857-999-6103

{guerster, belobaba, crawley, bcameron}@mit.edu

# SES S.A.  
Château de Betzdorf  
Betzdorf 6815  
Luxembourg

joel.grotz@ses.com

**Abstract** — In this paper we propose a Revenue Management framework for satcom operators and show with a proof-of-concept simulation that predicts a significant gain in revenues.

New satellite operators, highly variable demand for data, digital payloads, and new phased array technologies are likely to remake the current satcom landscape. One of the challenges operators old and new will face is how to manage demand and capacity. Airlines faced a similar situation with deregulation in the 1970s – their response with tiered pricing and seat inventory control to allocate capacity (known as Revenue Management), which may offer lessons for the satcom market.

The satcom industry shares many characteristics with the airline industry, such as inflexible capacity, low marginal sales cost, perishable inventory, heterogenous customers, and variable and uncertain demand. Generally, those characteristics favor the implementation of a Revenue Management system. However, the details of how Revenue Management can be used by satcom operators still need to be explored, which is the focus of this paper.

## TABLE OF CONTENTS

1. CHANGING LANDSCAPE OF SATELLITE OPERATORS .....	1
2. A PRIMER ON REVENUE MANAGEMENT .....	2
3. APPLICABILITY OF RM FOR SATCOM OPERATORS .....	3
4. COMPARISON TO PASSENGER AND AIR CARGO ...	5
5. PROPOSED RM SATCOM ARCHITECTURE.....	7
6. PROOF-OF-CONCEPT .....	8
7. CONCLUSION .....	12
ACKNOWLEDGMENT .....	13
REFERENCES.....	13
BIOGRAPHY.....	14

## 1. CHANGING LANDSCAPE OF SATELLITE OPERATORS

The satellite communication (satcom) world is divided into broadcasting and broadband services. Traditionally, the main business of satcom operator is the broadcasting of television content. Over the last decades, the consumers’ consumption is shifting from broadcasted content towards on-demand streaming, resulting in a shrinking satcom broadcasting sector. Therefore, the whole industry is moving towards the growing broadband sector, i.e., providing internet access to

airplanes, cruise and cargo ships, remote areas, backhauling of Wi-Fi and 4G/5G hotspots, and trunking.

In particular over the last few years, the transition towards broadband gained momentum with several companies entering the market. In particular, several Low Earth Orbit (LEO) constellations are proposed that provide throughput in the multiple Tbps range, with Telesat [1], OneWeb [2], and SpaceX [3] being most prominent. But there are also plans for smaller Medium Earth Orbit (MEO) and Geostationary Earth Orbit (GEO) constellation. For example, Viasat-3 will provide multiple 100 Mbps by 2020 [4]. SES will add with mPower seven satellite to their current O3b constellation [5] in 2021, and launch in the same year the GEO satellite SES-17 [6].

All of these new constellations fall into the class of high-throughput satellites (HTS). These new satellites are designed for different orbits and have different sizes, there is one key development that will change how satcom operators manage their satellites: *flexible payloads*. This flexibility is enabled by digital transparent processors and phased arrays, and it allows for the dynamic allocation of resources. Depending on the technical implementation, the degree of flexibility ranges from adjustable power only to a full control over frequency assignment, beam pointing, and beam shape.

While these new flexibilities offer great opportunities, they also come with challenges, in particular operational ones. Controlling the many degrees of freedom is no longer manually feasible and requires a Dynamic Resource Management (DRM) system (see Figure 1) [7].

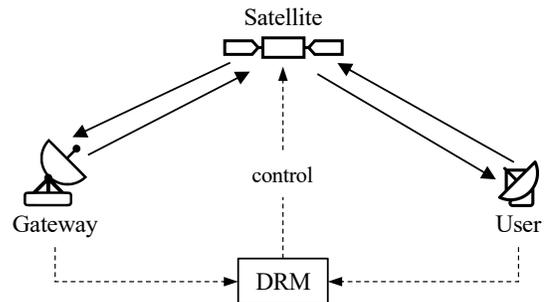


Figure 1: End-to-end system overview of satcom with Dynamic Resource Management (DRM)

The DRM receives input data from gateways and customers, generates an optimal control, and pushes it to the satellite. This in return affects the downlinks to gateways and customers, and hence closing the control loop. The DRM must be able to react dynamically to changes. This introduces new capabilities to regulate resources since the ground segment is now directly linked with the satellite payload control.

At the core of the DRM is an algorithm that handles the *resource allocation*. It determines how to match resources to a demand pattern while minimizing the satellites' resource consumption. The main resources of a satellite are its power and frequency spectrum. In fully flexible payloads, these are available as a pool and the DRM decides how to optimally split up the pool between the customers. The decisions fall into four categories: power, center frequency and amount of spectrum, beam pointing, and beam shape.

Most of the work on *resource allocation* is in the first category: power allocation. Author experiment with various algorithms techniques: Aravanis [8], Paris [9], He [10], Pachler [11], and Durand [12] use metaheuristics. Convex optimization is applied by Wang [13] and Deep Reinforcement Learning by Garau [14]. The frequency assignment problem is less studied with Paris [9] and Cocco [15] applying metaheuristics, Hu and Liu [16, 17] Deep Reinforcement Learning, and a Park [18] a binary search heuristic. Beam pointing and beam shape are addressed by Camino [19] with linear programming and by Kyrgiazos [20] with a self-organizing map. Pachler [21] developed two heuristic algorithm to solve the frequency assignment and beam pointing problem with particular focus on LEO constellations.

While the authors use different approaches and optimization techniques, they all arrive at the same conclusion: resource consumption is reduced by allocating resources more dynamically (enabled by flexible payloads). However, the current literature does *not* detail how the freed-up resources translate to a boost in revenues. Since the new digital payloads and phased arrays come with a cost premium, it is crucial for satcom operator to understand how they can leverage this technology. This opens up a research gap around the question:

**How can satcom operator monetize the flexibility of digital payloads and phased arrays?**

Addressing this question is the objective of this paper. We will build a conceptual framework drawing from Revenue Management (RM), connecting it to the DRM, and suggesting its applicability through simulation.

The paper is structured as follows. We provide a primer on RM in Section 2 and test its applicability for satcom in Section 3. A comparison with airline and air cargo in Section 4 leads us to the RM for satcom that we describe in Section 5 and validate in Section 6.

**2. A PRIMER ON REVENUE MANAGEMENT**

Revenue Management (RM) or also Yield Management is “*the process of allocating the right type of capacity to the right kind of customer at the right price so as to maximize revenues or yield*” [22]. More formally, RM is a conceptual framework for making sophisticated demand-management decisions in three categories: structural decisions about the selling format, pricing decisions, and quantity decisions [23].

American Airlines pioneered the development of many RM concepts and claimed in 1992 that RM increased their revenues by \$900M annually (in 2019 terms) [24, 25]. The underlying intuition is that empty seats are lost revenues. Any price above marginal cost could increase total revenues. At one extreme, selling all seats at just above marginal cost would result in a high proportion of seats filled (i.e., load factor), but not in a maximization of revenues. At the other extreme, selling only a few seats for very high prices is also not optimal, as many seats would remain empty [26]. Therefore, there is a trade-off to be made between the load factor and the average price (yield).

Further complexity is added when we consider the stochastic nature of demand. There are no-shows and late cancellations of bookings, which introduce uncertainty as to the actual number of empty seats remaining for sale. If customer segmentation is imperfect, lowering prices to attract customer from a segment with lower willingness-to-pay can lead to “buy-down” by customers with a higher willingness-to-pay. A RM system manages these decisions through optimization and forecasting. It uses historical data when available and learns continuously to make better decisions for future flights.

At the highest level, RM models can be divided into quantity- and price-based [23]. The difference is in the output of the optimization: either a maximum quantity that can be sold for a product with given price, or the price for a product given an estimated quantity. Historically, distributing changing prices was impractical and therefore airlines implemented quantity-based RM systems. However, with modern information technology this is no longer a limitation and specifically some new airlines have moved directly to price-based RM. In the following subsection we illustrate with a basic example the basic working principle of such deterministic *price-based* RM (adapted from Talluri [23]).

*Example of price-based RM working principles*

Let us consider two separate customer segments with the two demand elasticity functions  $d_1 = -2 \cdot p_1 + 260$  and  $d_2 = -p_2 + 150$  with  $p_1$  and  $p_2$  being the prices. We further denote the marginal revenues as

$$J(d) = \frac{\partial}{\partial d}(p \cdot d). \tag{1}$$

Therefore, we get  $J_1(d_1) = -d_1 + 130$  and  $J_2(d_2) = -2 \cdot d_2 + 150$ . We indicate the total revenues using  $\Pi$  and assume that capacity is limited to 100, i.e.,  $d_2 = 100 - d_1$ . The objective

is to find the optimal prices  $p_1^*$  and  $p_2^*$  so that  $\Pi$  is maximized. This is achieved when the marginal revenues  $J_1$  and  $J_2$  are equal. This problem can be solved by sweeping through all combinations of  $d_1$  and  $d_2$  as shown in Table 1.

**Table 1: Example for price-based RM**

$d_1$	$d_2$	$p_1$	$p_2$	$J_1(d_1)$	$J_2(d_2)$	$\Pi$
63	37	98.50	113.00	67	76	10,387
62	38	99.00	112.00	68	74	10,394
61	39	99.50	111.00	69	72	10,399
<b>60</b>	<b>40</b>	<b>100.00</b>	<b>110.00</b>	<b>70</b>	<b>70</b>	<b>10,400</b>
59	41	100.50	109.00	71	68	10,399
58	42	101.00	108.00	72	66	10,394
57	43	101.50	107.00	73	64	10,387

We find that the revenue maximizing prices are  $p_1^* = 100$  and  $p_2^* = 110$  (at this point the marginal revenues equal). The RM will then post these prices to the two customer segments (assuming this is possible). Given the price-elasticities, we expect a demand of  $d_1 = 60$  and  $d_2 = 40$ , clearing the capacity of 100. This example illustrates two segmentation scenarios:

One is where the selling takes places at a single time. In this scenario, the segmentation is based on customers' characteristics, e.g., the geographical location. With the example above: customer segment 1 might be from a more price-sensitive country than customer segment 2.

In the second scenario, the customer segments arrive during two separate time periods. For example, leisure air travel passengers tend to book earlier and business travelers later. This maps to the example's customer segment 1 and 2, respectively. The airline RM system will post a price that results in a demand of 60 for period 1 from leisure travelers with the expectation that it can charge more in period 2 for business travelers to clear the remaining 40. In a stochastic RM, these numbers are not point estimates but have an underlying probabilistic distribution [27].

In reality, RM often makes use of both segmentation strategies. Airlines used a Saturday-night-stay requirement to prevent business travelers from buying-down to cheaper fares even if they book early. And they have an advance purchase requirement for the cheaper fares, effectively raising prices the closer the booking gets to departure day. Very similar pricing schemes are exercised by hotels and rental car companies.

After the airline industry, hotels and rental cars were the early adapters of RM with many others following [28]. But what are the industry characteristics that favor the use of RM? And what does that mean for satcom operators? The subsequent text provides an answer to these questions.

#### *The six conditions for RM*

Various authors (e.g. Weatherford [29], Talluri [23] and Kimes [22]) derived sets of conditions that favor a RM. While

the sets are slightly different, the common six conditions in favor are:

- Capacity is inflexible
- Capacity costs are high compared to marginal sales cost
- Inventory is perishable
- Customers are heterogeneous and can be segmented
- Demand is variable and uncertain
- Organization has data and information system infrastructure

The goal of the following Section 3 is to test the suitability of RM for satcom operators against these conditions.

### **3. APPLICABILITY OF RM FOR SATCOM OPERATORS**

We use each of the six conditions that we introduced earlier to as a qualitative test of the potential applicability of RM for satcom operators. We describe them one by one below and summarize our findings at the end of the section.

#### *RM Condition 1: Capacity is inflexible*

The total capacity of communication satellites is defined during the design process and is fixed for the lifetime of the satellite (ranging between a few years and 15 years). In bent-pipe designs, the total capacity is fixed, as well as is the capacity distribution over the covered regions. With new flexible satellites, there is the option to reallocate capacity from one region to the other (the total capacity stays the same). Constellations with many smaller satellites have some flexibility to expand their capacity by launching new satellites or withdraw capacity by stopping replacement decommissioned satellites. Nevertheless, such changes are likely to take multiple years to result in significant capacity changes. As a result, capacity is considered inflexible which makes it essentially impossible for a satcom operator to match capacity to unpredictable short- and medium-term changes in demand. RM can reduce this mismatch by managing the demand side of the equation.

#### *RM Condition 2: Capacity costs are high compared to marginal sales cost*

The capacity costs (both in time and money) of launching a large communication satellite or a constellation of smaller satellites are substantial. They can range between many hundreds of million USD and a several billion USD. To assess the marginal sales costs, we consider two changes to the demand. First, demand changes for existing customers. Given the flexible nature of new satellites and an automated DRM, the marginal sales cost of this change is negligible. Second, adding a new customer: the marginal sale cost is mainly driven by the customer terminal, which ranges between multiple 100 USD for small GEO VSAT terminal and several hundred thousand USD for high-end 3-antenna

systems for MEO [30]. Emerging flat panel antennas have the potential to provide lower cost alternatives [31]. Low marginal sale costs allow a RM to change demand more frequently and therefore be more reactive. But, as we illustrated, terminal costs are not negligible and hence limit the frequency and extent of changes. This observation leads us to the separation into more frequent *contract add-ons* (changes to existing customers) and less frequent *long-term allotments* (adding a new customers). As we will see in Section 4, this characteristic has a strong analogy with air cargo.

*RM Condition 3: Inventory is perishable*

The inventory of satcom operators is the power and frequency spectrum available on their satellites. Perishability means that when a resource is not used at the moment, the unused resources cannot be stored and used at a later time. The resource is spoiled and has an opportunity cost. The perishability property is true for spectrum, but not necessarily for power: it can be stored in on-board batteries. The storage capabilities are limited and depend on design, orbit, and demand pattern. The partial imperishability of power introduces a time-dependency into the RM formulation on capacity level. This is not a limitation per se but adds complexity. A reasonable first approximation is to assume a constant average power limit per orbit and consider power to be perishable.

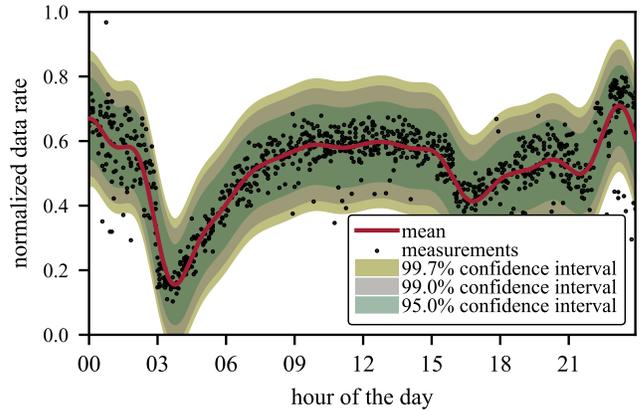
*RM Condition 4: Customer are heterogeneous and can be segmented*

Satcom operators have a variety of customers. An example segmentation is along the following three dimensions. One is the type and size of the customer: a single terminal end-customer, a multi-location end-customer, or a wholesaler. Another can be based on geographical attributes, such as the country or the location of the terminals: land, sea, or air. Third, the gateways can be owned and operated by the satcom operator or by the customers themselves. All of these three dimensions are directly observable by the satcom operator and hence allow for the potential of effective segmentation. A successful segmentation approach is stable over time and has a substantial, homogenous customer base within each segment [32]. The more heterogenous the segments are between themselves, the better the RM can differentiate between them.

*RM Condition 5: Demand has variability and uncertainty*

Variability are changes over time. They can occur on different timescales (hours, days, months, ...) and can be periodic, repeating, or single events. Depending on the nature of the variation, a prediction of it has higher or lower uncertainty. In addition, there is an uncertainty about future demand: it depends on how the market evolves, how competitors behave, how technology develops, etc. To understand the variability and uncertainty of demand for satcom, we observed and analyzed data from SES's customer [33]. We detected that depending on the customer, the variability and uncertainty is considerably different. The

relative uncertainty becomes smaller the more demand is aggregated. We also note, that the most distinct periodic pattern is diurnal with differences in day and night usage exceeding 50% in many cases, see Figure 2 for an example. We fitted a Gaussian Process to one-month worth of actual measurements. We randomly sampled one thousand points and used a kernel with a sum of a radial basis function and white noise. The drop in night usage is closely correlated with the local time-zone and therefore longitude. From the view of the satellite, this results in an additional geographical variation along the longitude. RM can exploit these variations and uncertainties by statistically multiplex customers to better utilize the satellite.



**Figure 2: Normalized 24h data rate variation**

*RM Condition 6: Organization has data and information system infrastructure*

While the first five conditions are at a market level, the last condition is organizational. We cannot make a general assessment for all satcom operators, but we will use the example of SES S.A. The company plans to launch the highly flexible MEO constellation mPower in 2021. To operate and control the satellites' many degrees of freedom, a DRM is under development. A significant part of this development effort is to integrate data streams from ground segments and satellites into a common accessible place. The optimization engine of this tool then uses this data to automatically generate a new configuration that is automatically commanded to the satellites. The deployed data and information system infrastructure can be used as foundation for the development of RM for satcom. This synergy helps to shorten development time and reduces cost and risk.

*Summary of findings*

We examined the satellite industry against the six conditions that favor the implementation of RM. We conclude that the marginal sales cost for adding a new customer are not negligible and therefore we separated the changes to demand into contract add-ons and long-term allotments. We conclude that all conditions are met, which supports our proposal of RM for satcom operators. In the following Section 4 we identify the similarities and differences of satcom operators with passenger and cargo airlines.

#### 4. COMPARISON TO PASSENGER AND AIR CARGO

The objective of this Section is to compare a proposed satcom RM to existing practice of RM. The two industries we use for comparison are passenger airlines and air cargo. RM for passenger airlines has an exhaustive literature and is the pioneering success story of RM. We will illustrate that air cargo has many similarities with satcom from which we can draw insights. In particular, we look at four comparisons: characteristics of capacity, characteristics of demand, overbooking, and current RM practices. We discuss each of them in the following and summarize them in Table 2.

##### *Capacity characteristics*

The capacity for passenger airlines is one-dimensional and measured in seats. Since an airplane is assigned to a flight segment well ahead of time, the number of seats in each cabin is fixed and known. There are some exceptions: e.g. the movable curtain between business and economy class for mostly European airlines [28]. This strategy allows for changing the capacity of each cabin more flexibly. However, generally speaking, capacity for airlines is one-dimensional, discrete, fixed and known.

One of the significant differences between passenger airline and air cargo is that capacity for cargo is multi-dimensional [34, 35]. The two main properties of a shipment are its continuous mass and volume. A standard density is sometimes used to collapse them into one dimension for complexity reduction [36]. Each flight has a variable and uncertain available capacity for cargo. For example, the loaded fuel is a function of the weather conditions and therefore influence the available capacity. For combined passenger-cargo flights, the available cargo capacity can depend on the number of passengers and their bags.

The capacity for satcom can be measured in available power and frequency spectrum. In contrast to cargo’s mass and

volume, power and spectrum can be traded off to achieve the same demanded data rate. With more spectrum, the spectral efficiency can be decreased while still achieving the same data rate. A lower spectral efficiency requires less power from the link, and hence power can be reduced. Since power is less perishable than spectrum, priority should be given to fill-up the available spectrum first and adapt to changes in the demanded data rate by changing power. With that, we can formulate a continuous one-dimensional capacity in terms of power. Since weather is difficult to predict, it introduces uncertainty. In air cargo it affects the loaded fuel, and for satcom it affects the quality of the link. For example, during rain, the power has to be increased to keep the link’s data rate constant and therefore the available capacity is reduced. Uncertainties around legal matters of frequency rights can create additional medium-term variation of the available capacity. Furthermore, the capacity varies over the long-term as the satellite ages. The performance of solar cells and batteries decline over time and impact the pool of available power.

##### *Demand characteristics*

Passenger airline demand is itinerary-based, i.e. a seat for an origin-destination itinerary on a specific flight or flights. Larger group bookings and contingent reservations are rather the exception. The bookings are for a single departure of the flight and therefore short-term.

In air cargo, the demand is similar to the capacity – multi-dimensional with mass and volume. But compared to the passenger airlines, the demand is generally service oriented [34]. This means that the shipment is not bound to a specific flight or route, but rather has a time/date by which the shipper guarantees delivery at the destination. The most significant difference is that air cargo shippers typically have a few large customer who reserve a majority of the available capacity [35]. Since these contracts are often renegotiated

**Table 2: Comparison of the characteristics between satcom operators and airline and cargo airlines**

	<b>Passenger airlines</b>	<b>Air cargo</b>	<b>Satcom operators</b>
<b>Capacity characteristics</b>	<ul style="list-style-type: none"> <li>• discrete</li> <li>• 1D: seats</li> <li>• fixed and known</li> </ul>	<ul style="list-style-type: none"> <li>• continuous</li> <li>• 2D: mass, volume</li> <li>• variable and uncertain</li> </ul>	<ul style="list-style-type: none"> <li>• continuous</li> <li>• 1D: power</li> <li>• variable and uncertain</li> </ul>
<b>Demand characteristics</b>	<ul style="list-style-type: none"> <li>• short-term bookings</li> <li>• 1D: seats</li> </ul>	<ul style="list-style-type: none"> <li>• medium-term allotments</li> <li>• short-term bookings</li> <li>• 2D: mass, volume</li> </ul>	<ul style="list-style-type: none"> <li>• long-term allotments</li> <li>• contract add-ons</li> <li>• 1D: data rate (requires a model to map data rate to power)</li> </ul>
<b>Overbooking based on</b>	<ul style="list-style-type: none"> <li>• no-show rate</li> <li>• late cancellations</li> </ul>	<ul style="list-style-type: none"> <li>• no-show rate of short-term bookings</li> <li>• uncertainty in capacity usage of medium-term allotments</li> <li>• for combined passenger cargo flights: number of passenger and bags</li> </ul>	<ul style="list-style-type: none"> <li>• no-show rate and usage of contract add-ons</li> <li>• variation and uncertainty in capacity usage of long-term allotments</li> </ul>
<b>RM practice and decisions</b>	Quantity-based RM with transition to price-based RM; decisions: <ul style="list-style-type: none"> <li>• quantity of each booking class</li> </ul>	Quantity-based RM; decisions: <ul style="list-style-type: none"> <li>• quantity of short-term bookings</li> <li>• ratio of short-term bookings to medium-term allotments</li> <li>• quantity for medium-term allotments</li> </ul>	Price-based RM; decisions: <ul style="list-style-type: none"> <li>• price for contract add-ons</li> <li>• price for long-term allotments</li> </ul>

twice per year [37], they are commonly referred to as medium-term allotments. In addition to this demand, there are also short-term bookings for not recurring shipments.

Satcom operators share many similarities with air cargo shippers when it comes to demand. The bookings are service bound through Service Level Agreements (SLAs) and there are a few large customers who reserve a majority of the available capacity. Current SLAs have three main parameters: Committed Information Rate (CIR), uptime of the link (= availability), and a monthly price for the CIR. The satcom operator promises to provide the CIR for a defined uptime (usually expressed in percentage). If the SLA is violated, the customer receives a credit. Under this contract, provisioning of a data rate above CIR, is not tied to additional revenues. Since these SLAs are typically ranging from one to five years [30], we refer to them as *long-term allotments*. This supports our observation from Section 3: marginal sales cost of adding a new customer are considerable and therefore longer-term contracts are favored. However, we also found that marginal sales costs for changing the data rate for existing customer is negligible. This opens up a new opportunity for *contract add-ons* similar to short-term bookings for passenger and cargo airlines. We elaborate on this idea further in Section 5. A uniqueness of the satcom RM is that demand and capacity have different units. The demand is a data rate and the capacity is power. The relationship between these two units is non-linear and time dependent for Non-geostationary satellites (NGSO) or moving terminals. We therefore need a *model* that translates between the two units.

### *Overbooking*

Whenever either capacity or demand is variable or uncertain, overbooking aims to minimize the mismatch between capacity and demand. For passenger airlines the uncertainty is on the demand side with no-shows and late cancellations of passengers. No-show rates have historically averaged at around 10-15% and therefore overbooking has accounted for around half of the 4-6% revenue gain of RM [38]. The other half is attributed to fare class optimization.

In air cargo, overbooking has two additional sources. One rises from the actual usage of medium-term allotment. There is uncertainty in how customers use their reserved allotments. The second one is uncertainty in the available capacity through weather conditions, or for combined flights, the passengers' no-show rate and number of bags.

Satcom can also have no-shows of contract add-ons. In addition, the actual usage of the add-ons might differ from the expected profile. Long-term allotments offer the most significant opportunity for overbooking, however. We discussed in Section 3 that demand has a strong diurnal variation. The specific form of this variation has an underlying uncertainty. Overbooking can make use of these variations and uncertainties to bring demand closer to capacity.

### *RM practice and decisions*

As outlined in Section 2, it was historically difficult to change prices dynamically and therefore airlines implemented quantity-based RM systems. The decisions that RM for airlines takes is the quantity for each bookings class. A booking system then closes the lower, cheaper booking classes when this quantity is reached and only the higher, more expensive booking classes remain open. This generates the impression to the customer that the airline prices dynamically, when in fact the quantity-based RM only does this indirectly.

We found mostly quantity-based RM for air cargo [34-37]. The cargo RM system makes decision at different time scales. The quantity of allowed short-term bookings has to be determined to match demand to capacity without violating medium-term allotments. The ratio between short-term bookings and medium-term allotments is another parameter that becomes important whenever contracts for allotments expire or new allotments are negotiated. Along with this, the quantity allocated to each medium-term allotment is of importance.

For satcom, again, many similarities can be drawn from air cargo. Since distributing prices is no longer a limitation in today's internet age, a price-based RM approach is the best fit. The decisions are then to set the price for contract add-ons based on an expected quantity that bring demand close to capacity. Whenever the long-term allotments are negotiated, the RM system informs this process with prices. By posting these prices, the ratio between contract add-ons and long-term allotments is determined by the market.

### *Summary of similarities and areas of further attention*

We contextualized the satcom RM with existing practices for passenger airline and air cargo. Many similarities allow us to build a satcom RM system on an exhaustive set of existing literature. Some of the similarities are drawn from passenger airlines and most of them from air cargo. However, we also identified three areas that need further attention. These are:

- Diurnal variations of *long-term allotments* offer a significant opportunity for satcom operator. Since they have an underlying uncertainty, they have to be *estimated*.
- Demand has a different unit than capacity and therefore a *model* is needed for translation.
- Satcom operators do not have a parametrized product menu for *contract add-ons* yet.

In the subsequent Section we describe more specifically the potential for satcom RM with particular focus on these three areas.

## 5. PROPOSED RM SATCOM FRAMEWORK

Talluri [23] decomposes RM into four layers: data input, estimation and forecasting, optimization, and control. Data is collected from different locations, conditioned if needed, and stored in one place. From this an estimator is built that also can be used to forecast. Once an optimization is triggered, the estimates and forecasts are used to find a set of optimal policies. The final control step makes decisions based on the optimal policies and manages the day-to-day booking transactions.

The objective of this Section is to specify the building blocks of these four layers that then define a possible satcom RM system. Figure 3 provides an overview of the building blocks that we describe in detail below. There are three key blocks that are specific to satcom: available capacity forecasting, parametrized contract add-on menu, and model.

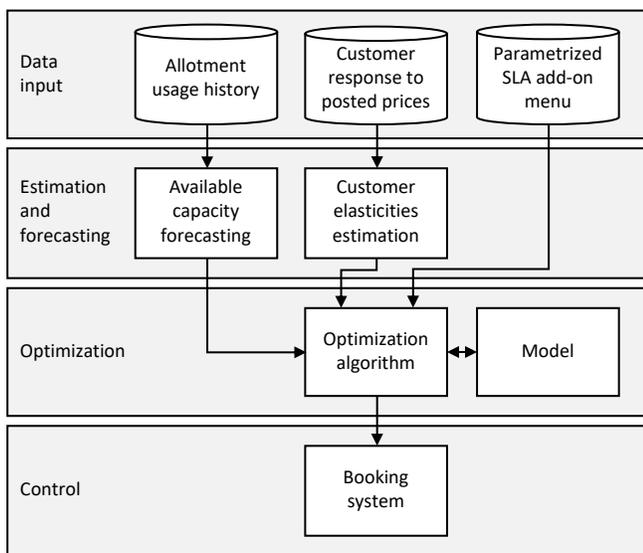


Figure 3: Overview of the RM framework for satcom

### Data input

In our framework, we show three databases that store information. In reality they have interfaces to collect data from various systems inside and outside the company. Before data is processed by the estimator and forecaster, it has to be conditioned. The three databases are:

**Allotment usage history.** This database is directly fed by the modems that measure the data rate at the terminal of long-term allotment customers. The characteristics of the data streams vary by modem vendor, so all data need to be brought into a common form before storing. Adding information about special events or anomalies helps to build a more accurate estimator.

**Customer response to posted prices.** In the beginning, this database will only have customers' response to long-term allotment pricing. For contract add-ons it will be empty and expert knowledge has to be inserted to determine the first set of prices. But over time the database will grow as customer respond to the pricing of contract add-ons.

**Parametrized SLA add-on menu.** Basically, all products that satcom operator currently sell are long-term allotment SLAs. They are not formally parameterized but rather written documents and often individualized [30]. For a RM to work, the products need to be parameterized. This is particularly important for contract add-ons as the idea behind these products is to sell them more frequently in an automated matter. The thorough development of a product menu is outside of the scope of this paper, nevertheless we provide an example add-on menu in Section 6 as part of our proof-of-concept.

### Estimation and forecasting

Estimation is descriptive, it finds the parameters of a model to best fit observed data. Forecasting is predictive, it uses this model to calculate future unobserved values [23]. The input is conditioned data and the output is a probability distribution or a point estimate. We have two building blocks in this layer:

**Available capacity forecasting.** This forecaster is directly fed by the allotment usage history database. The goal of the forecaster is to provide the optimization algorithm with a probability density function (PDF) of available capacity for any given time. As this is one of the key building blocks of the satcom RM, we will provide an example in Section 6.

**Customer elasticities estimation.** Having an understanding of how customer will respond to pricing in different segments is key to every RM system. Fed by the customer response database, this building block's objective is to estimate various elasticities. Extensive work has been done to build demand elasticities estimators and we argue that they can be well integrated into a potential satcom RM system.

### Optimization

The optimization layer consists of two blocks: the algorithm and the model.

**Algorithms.** The inputs of the algorithm are the forecast of the available capacity, the demand elasticity estimates, and the parameterized contract add-on menu. The algorithm then optimizes within these given constraints to produce a set of optimal prices. These prices are then passed onto the next layer. As the inputs have different units, a model needs to be called to translate them into one common unit as a function of the time. Since there is a vast literature on algorithms available, we argue that the blocks can be tweaked to fit existing algorithms (as we illustrated in the proof-of-concept Section 6).

**Model.** As a unique quality of the satcom RM, capacity and demand have different units. The relationship is non-linear and is a function of time (at least one of the satellite or the terminal is moving). The details of this model are outside of the scope of this paper, but we develop an approximation in Section 6.

## Control with a booking system

The control layer consists of the booking system that makes decisions based on the optimal prices from the optimization algorithm. It pushes the prices to customer, records their response, and accepts and rejects bookings. There is a feedback loop to the available capacity forecaster and to the customer response database.

## 6. PROOF-OF-CONCEPT

The purpose of this Section is to provide a proof-of-concept of the framework for a satcom RM. We will discuss one possible realization of such a system to illustrate the working principles and the potential gains. Since HTS with flexible satellites are still to be launched, we rely on simulation. We first describe the setup and the assumptions made. Then, we will go into detail on the three key blocks model, available capacity forecaster, and parametrized SLA add-on menu. For the latter we assume a single product that we call *spot instance*. We use an optimization technique based on marginal revenues similar to the one introduced in Section 2. Finally, results show the revenue gains of implementing RM.

### Setup of the simulation and assumptions

We consider a GEO satellite located at 0° latitude and 0° longitude with an altitude of 35,629 km (see Figure 4). We focus on the forward downlink from satellite to customer. Each customer has its own beam with a half-cone angle of 0.7°. The satellite serves one customer from each of the five segments: aviation, backhauling, maritime, trunking, and vsat (very small aperture terminal).

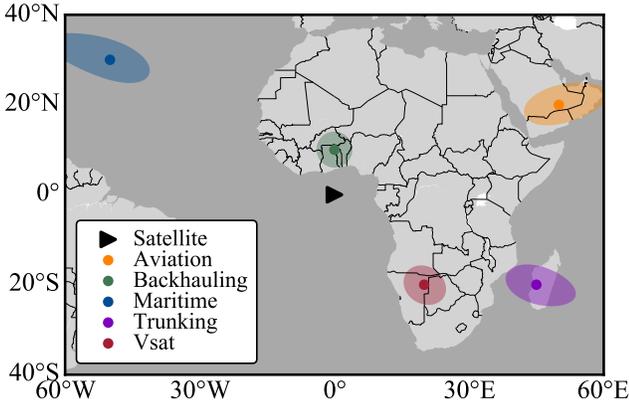


Figure 4: Overview map of the location of the five customers and the GEO satellite.

Each customer segment has a different characteristic which we describe in the following:

**Aviation** services provide broadband connectivity to commercial or private aircrafts. The end-customers are the passengers on board. All passengers are multiplexed and perceived as a single customer by the satellite.

**Backhauling** services connect a subnetwork with the fiber backbone of the Internet. Often the subnetwork is a cellphone tower or Wi-Fi hotspot in regions with no or low speed terrestrial connection. Backhauling is often on larger scales and is especially considered a viable solution for developing countries.

**Maritime** services are similar to aviation. Due to slower movements, fewer dynamics are added to the communication links. For large cruise ships, thousands of end-customers are multiplexed. The cruise ship cooperation often acts as intermediary between satellite operator and end-customer.

**Trunking** services are similar to Backhauling with the focus on using satellites for excess demand in terrestrial networks or for contingency scenarios. The traffic volume is often considerable, and less variable and uncertain as many end-customers are multiplexed.

**Vsat** services address private end-customer and business solutions with light and variable traffic demands. The satellite operator often directly sells their services to the end-customer and provides connectivity to the Internet backbone.

Based on these descriptions we parametrize five representative customers as shown in Table 3 and illustrated in Figure 5. We group them into three categories: technical, demand, and SLA parameters. For the technical, we list the longitudes (*lon*), latitudes (*lat*), the diameter of the terminal dish *D*, the receiver figure of merit *G/T*, and a desired spectral efficiency  $\Gamma$  of the link.

The mean data rate  $\mu(t)$  of the 24h demand pattern is modelled by a cosine with three parameters and is scaled by the CIR. Eq. (2) describes the relationship

$$\mu(t) = CIR \cdot \left( \mu_b + v \cdot \cos\left(\frac{t - \Delta t}{12} - \pi\right) \right) \quad (2)$$

with *t* being the hour of the day (in Coordinated Universal Time UTC),  $\mu_b$  the normalized base mean, *v* the variation of

Table 3: Assumptions about the customers input parameters.

	Technical parameters					Demand parameters				SLA parameters			
	<i>Lon</i> [°]	<i>Lat</i> [°]	<i>D</i> [m]	<i>G/T</i> [dBK]	$\Gamma$ [-]	$\mu_b$ [-]	<i>v</i> [-]	$\Delta t$ [h]	$\sigma_b$ [-]	<i>CIR</i> [Mbps]	$p_{CIR}$ [\$/month/Mbps]	$\epsilon_{CIR}$ [-]	<i>A</i> [-]
<b>Aviation</b>	50	20	0.6	15	1	0.4	0.30	7.33	0.3	70	300	0.5	0.99
<b>Backhauling</b>	0	10	1.2	21	2	0.5	0.30	4.00	0.2	150	100	0.9	0.99
<b>Maritime</b>	-50	30	1.2	21	2	0.4	0.20	0.67	0.4	100	200	0.7	0.99
<b>Trunking</b>	45	-20	2.4	27	3	0.7	0.05	7.00	0.1	250	100	0.8	0.99
<b>Vsat</b>	20	-20	1.2	21	2	0.3	0.20	5.33	0.5	50	200	0.7	0.99

the cosine around  $\mu_b$ , and  $\Delta t$  the location of the minimum usage, i.e. the hour of the night drop with respect to UTC. The uncertainty around the mean is assumed to be proportional to the mean with a base sigma  $\sigma_b$ :

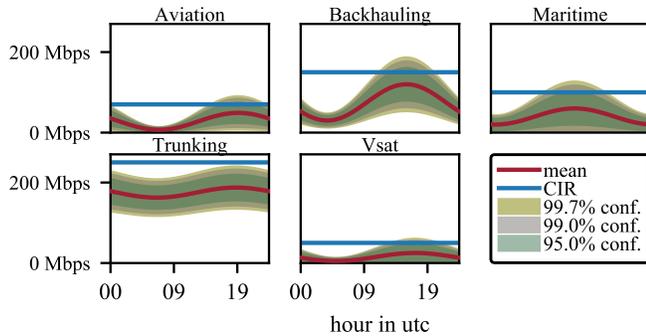
$$\sigma(t) = \sigma_b \cdot R(t) \quad (3)$$

Furthermore, we assume the demand has a normal distribution and therefore we define a stochastic process as a collection of independent random normals:

$$\{\mathbf{X}_t\}_{t \in T} \quad \text{with} \quad \mathbf{X}_t \sim N(\mu(t), \sigma(t)^2) \quad (4)$$

with  $\mathbf{X}_t$  being a normal random variable and  $T$  being the 24 hours of the day. For the computations, we discretize  $T$  into 15 minutes. The demanded data rate  $R(t)$  is then a sample from  $\mathbf{X}_t$ .

Applying the parameters listed Table 3 to Eqs. (2) - (4), we get Figure 5, which displays the demand pattern of the five customers introduced. All customers are modelled with a drop in night usage over a seven hour interval. Trunking is the largest customers with a CIR of 250 Mbps, with a relatively small variation  $\nu$  and uncertainty  $\sigma$ . In contrast, Vsat has the lowest CIR with 50 Mbps and a higher uncertainty around the mean. Aviation and backhauling are in-between with aviation being more uncertain but uses on average 40% of the CIR while backhauling uses 50% ( $\mu_b = 0.5$ ). Maritime is modelled with a higher uncertainty due to the more bursty characteristic of the customer segment.



**Figure 5: Five customers modelled over a 24 h period by cosines. The confidence intervals are a function of the mean (see Eq. (3)).**

### Model

The relationship between the demanded data rate  $R$  and the power capacity  $C$  of the satellite is defined by the link budget [39]. Since we consider a GEO satellite, the geometry is not time-dependent and simplifies the model. For analytical simplicity, we use the Shannon limit  $R = B \cdot \log_2(1 + C/N)$  so we can write for customer  $i$  in linear form

$$R_i(t) = B_i \cdot \log_2 \left( 1 + \frac{C_i(t) \cdot G_{Tx} \cdot G/T_i}{L \cdot k_B \cdot B_i} \right) \quad (5)$$

with  $L$  being the free space loss that we assume here to be constant with 209.5 dBi. The gain of the satellite is  $G_{Tx}$  and

assumed to be 51 dBi. The variable  $k_B$  is the Boltzmann constant with -228.6 dBK. Solving for  $C_i(t)$  gives us Eq. (6) in which we define  $Z_i$  to be the first term.

$$C_i(t) = \frac{L \cdot k_B \cdot B_i}{\underbrace{G_{Tx} \cdot G/T_i}_{Z_i}} \cdot \left( 2^{\frac{R_i(t)}{B_i}} - 1 \right) \quad (6)$$

The exponent in the parentheses introduces a non-linearity that together with the customer specific properties  $G/T_i$  and  $B_i$  lead to significant differences in the capacity cost per data rate  $C(t)/R(t)$ . Table 4 illustrates that with using the  $R(t) = CIR$ .

**Table 4: different capacity costs of the five customers**

	Specific capacity cost [mW/Mbps]	Cost increase relative to trunking
<b>Aviation</b>	5.7	682%
<b>Backhauling</b>	2.9	348%
<b>Maritime</b>	4.8	572%
<b>Trunking</b>	0.8	100%
<b>Vsat</b>	4.8	572%

The first column shows the specific capacity cost in  $mW/Mbps$  and the second column shows the cost increase relative to trunking which is the “cheapest”. Given our assumptions, backhauling is 3x, maritime and vsat are almost 6x and aviation is almost 7x times more expensive to serve than trunking. The largest driver for this is the difference in their  $G/T$ . The optimization will trade-off the higher cost with the higher willingness to pay of some segments.

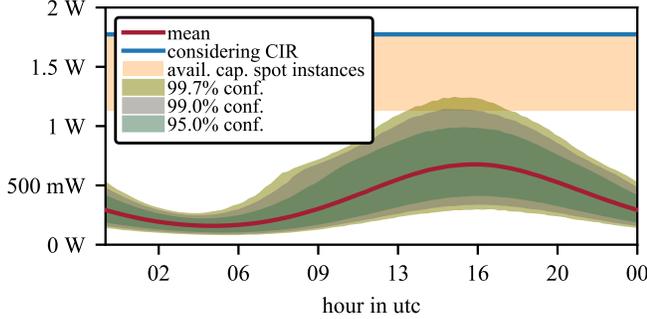
### Available capacity forecaster

The available capacity forecaster determines how much capacity is available for short-term contract add-ons. It sums up the required capacity from Eq. (6) for each customer  $i$  and subtracts it from the maximum capacity  $C_{max}$ . In the deterministic case, we formally have

$$C_{avail}(t) = C_{max} - \sum_{i=1}^n C_i(t) \quad (7)$$

with  $n$  being 5 in this example. We set  $C_{max}$  to the power required if capacity is allocated based on CIR, i.e.,  $C_{max} = 1.77 W$ . As we work with stochastic processes, we sample 50,000 times  $R_i(t)$  from the demand  $\mathbf{X}_{i,t}$  and get a distribution of possible  $C_{avail}(t)$  (assuming that demand is not correlated). Since the distribution is no longer a normal random variable, we work with an empirical distribution function for each  $t$ . We denote the resulting random variable as  $\mathbf{Y}_t$ . Figure 6 shows the resulting process  $\{\mathbf{Y}_t\}_{t \in T}$  including its confidence intervals. The blue line is the capacity that would have to be reserve if allocating for CIR, i.e., without having flexibility in the payloads to follow the demand pattern and reallocated capacity.

Due to the non-linearities, the uncertainty is no longer symmetric around the mean, particularly for higher data rates at Hour 15. The daily pattern follows that of the input demand data with a nigh-drop around Hour 5. While the 99 % upper confidence interval of the demand was touching the CIR line for the high usage period (see Figure 5), we can see from Figure 6 the gains of statistical multiplexing of customers. There is capacity available that can be used, even in the high usage period.



**Figure 6: Plot of the used capacity modelled by the stochastic process  $\{Y_t\}_{t \in T}$  after sampling  $X_{i,t}$  50,000 times for each  $i$  and  $t$ . The blue line illustrated the maximum capacity which is equal to an allocation based on CIRs and not following the demand.**

To calculate the available capacity in this stochastic case, we can rewrite Eq. (7) with the promised availability  $A$  from Table 3 to

$$\mathbb{P}(Y_t < C_{max} - C_{avail}(t)) \geq A. \quad (8)$$

We solve this equation for  $C_{avail}(t)$  by inverting the empirical cumulative distribution function of  $Y_t$ . This gives us how much of the capacity is available for reallocation at each time  $t$  while ensuring the promised availability  $A$  of the allotment SLA is met.

#### Parametrized SLA add-on menu: spot instance

For the purpose of this example, we define one product that makes use of the available capacity. We call it *spot instance* (in analogy to the alike Amazon Web Service product [40]). It is an add-on to the existing SLAs and is sold to the *same* customers. It works with the customers' price elasticity to generate additional demand. The price is below the long-term allotments, but the operator can quit the service any time. This achieves a win-win situation for both sides. The customer receives additional data rate for a lower price and the operator can safely overbook the satellite: if the usage of the long-term allotment changes, the operator can adjust to these changes by quitting spot instances. In contrast, if the operator overbooked with long-term allotments and usage behavior changes, adjustments can only be made after the long-term SLAs expire or are renegotiated.

The parameters of the *spot instance* are similar to the long-term allotment SLA. The product has a data rate and an associated price. We assume perfect customer segmentation

is possible and therefore the parameters are different for each customer  $i$ :  $R_{spot,i}$  and  $p_{spot,i}$ .

For further simplification, we make the conservative assumption that all customers use their full  $R_{spot,i}$  throughout the 24 h window. With that, the most congested point is during peak hour of the existing SLA allotments. We get the available capacity by finding the minimum of all  $C_{avail}(t)$ . Formally we have:

$$C_{avail} = \min(C_{avail}(t) \forall t \in T) \quad (9)$$

which results in  $C_{avail} = 0.63 W$  (see orange area in Figure 6) in our example. Note, that there is still the white area remaining between the 99% confidence interval line and the orange area. Hence, with this assumption we only consider approximately half of the theoretical free capacity as available – the rest is the undulating white space shown below the orange area and above the green used capacity.

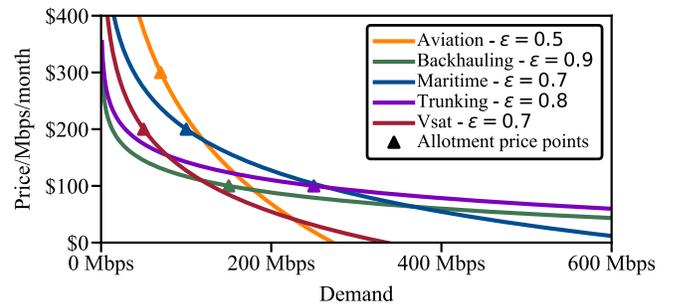
For the customer demand elasticity, we assume a log-linear (exponential) behavior that goes through the price point defined by the long-term allotment. With  $R_{spot,i}$  being the additional data rate contracted, we get the price  $p_{spot,i}$  by

$$p_{spot,i} = \frac{1}{b_i} \cdot \ln\left(\frac{a_i}{CIR_i + R_{spot,i}}\right). \quad (10)$$

The elasticity  $\epsilon$  for the log-linear demand is defined by Eq. (11) [23].

$$\epsilon_{CIR,i} = p_{CIR,i} \cdot b_i \cdot e^{-1} \quad (11)$$

We calculate the parameter  $b_i$  by using the elasticity defined for  $p_{CIR,i}$  from Table 3. With the values for  $p_{CIR,i}$  and  $R_{spot,i} = 0$  we obtain the parameter  $a_i$ . The resulting relationships are plotted in Figure 7. Aviation is modelled as the most price-insensitive customer segment, following by maritime and vsat, trunking, and backhauling. The upwards pointing triangle is the assumed price point of the long-term SLAs.



**Figure 7: Lin-log demand functions for the five customers. The triangle is the price point of the allotments for which the elasticity  $\epsilon$  is defined.**

Furthermore, with the known parameters  $a_i$  and  $b_i$ , the additional revenue  $\Pi_{spot,i}$  is obtained by multiplying Eq. (10) by  $R_{spot,i}$ :

$$\Pi_{spot,i} = \frac{R_{spot,i}}{b_i} \cdot \ln\left(\frac{a_i}{CIR_i + R_{spot,i}}\right) \quad (12)$$

The total revenues  $\Pi_{tot}$  are then defined as

$$\Pi_{tot} = \Pi_{CIR} + \Pi_{spot} \quad (13)$$

with  $\Pi_{CIR} = \sum_{i=1}^n p_{i,CIR} \cdot CIR_i$  and  $\Pi_{spot} = \sum_{i=1}^n \Pi_{spot,i}$  for customer  $i$ .

### Revenue optimization

For revenue maximization we use a gradient-based greedy algorithm similar to the one introduced in Section 2. We analytically derive the gradient  $\partial\Pi_{spot,i}/\partial C_{spot,i}$  by inserting Eq. (5) into Eq. (12) (thus assuming the *spot instance* is provided through a new beam to the same terminal). For the purpose of clearer notation, we omit the customer index  $spot, i$ . With the defined variable  $Z = \frac{L \cdot k_B \cdot B}{G_{Tx} \cdot G/T}$  from Eq. (6) we get

$$\frac{\partial\Pi}{\partial C} = \frac{\partial}{\partial C} \left( \frac{B \log_2\left(\frac{C}{Z} + 1\right)}{b} \ln\left(\frac{a}{CIR + B \log_2\left(\frac{C}{Z} + 1\right)}\right) \right) \quad (14)$$

With the substitutions

$$D_1 = \frac{C}{Z} + 1 \quad \text{and} \quad D_2 = B \cdot \frac{\ln(D_1)}{\ln(2)} + CIR \quad (15)$$

the final gradient reads

$$\frac{\partial\Pi}{\partial C} = \frac{B \cdot \ln\left(\frac{a}{D_2}\right)}{b \cdot Z \cdot \ln(2) \cdot D_1} - \frac{B^2 \cdot \ln(D_1)}{b \cdot Z \cdot \ln^2(2) \cdot D_1 \cdot D_2} \quad (16)$$

We calculate these gradients for every customer  $i$  and then follow the procedure outlined in Algorithm 1.

---

### Algorithm 1: Gradient-based revenue optimizer

---

**Input:**  $J_i(C_i) = \partial\Pi_i/\partial C_i$   
**Input:** step size  $\Delta C$   
**Input:** available capacity  $C_{avail}$   
**Output:** optimal allocation  $C_i^*$

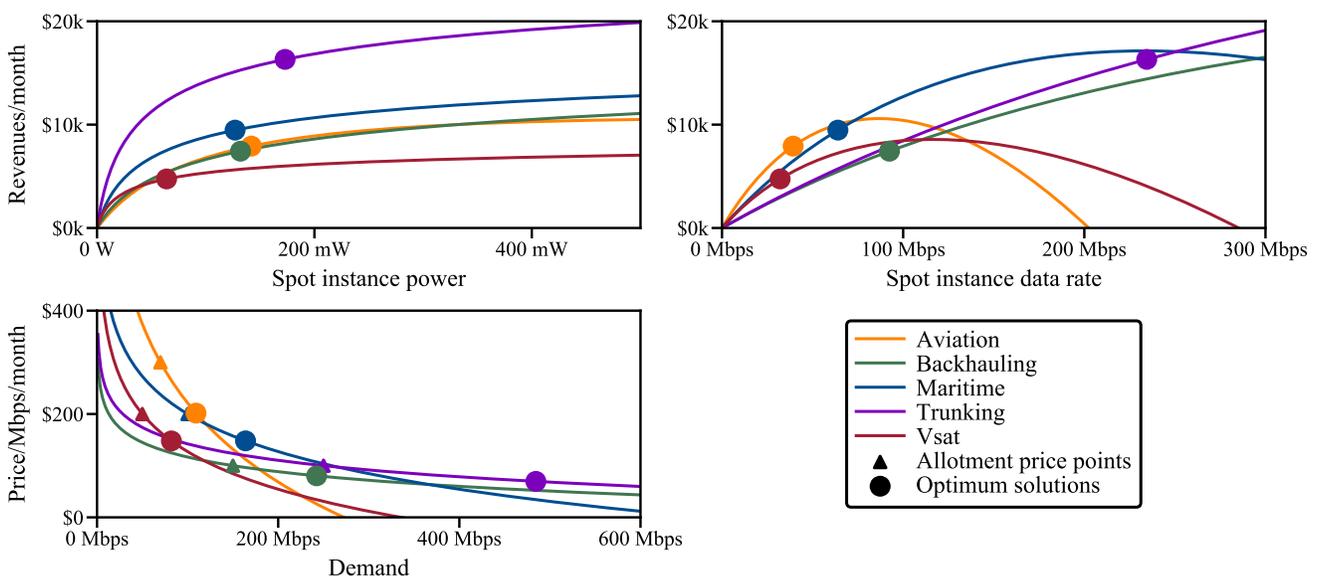
- 1: Initialize  $J_i(C_i)$  with  $C_i = 0$
- 2: **while**  $\sum_{i=1}^n C_i < C_{avail}$  **do** // repeat until capacity limit reached
- 3:   **if** any  $J_i(C_i) > 0$  **then** // check if any gradient positive
- 4:      $C_{argmax_i(J_i)} += \Delta C$  // incr. capacity for steepest grad.
- 5:   **else**
- 6:     **stop** // stop when all gradients negative
- 7:   **end if**
- 8: **end while**

---

The gradients are initialized with a capacity  $C_i = 0$ . Then, iteratively the capacity is increased by  $\Delta C$  for the customer  $i$  with the steepest gradient, i.e., the highest marginal revenue. For the step size we choose  $\Delta C = 0.001 W$ . The algorithm terminates if either all gradients are negative (selling more capacity would reduce revenues) or the capacity limit  $C_{avail}$  is reached (capacity cleared). The result is the revenue maximizing allocation  $C_{spot,i}^*$  for each customer  $i$ . With Eq. (5), we transform this capacity back into  $R_{spot,i}^*$  and with the demand function in Eq. (10) we get the optimal prices  $p_{spot,i}^*$ . Since the marginal revenue function  $\partial\Pi_{spot,i}/\partial C_{spot,i}$  is monotonically decreasing, the found solution is ensured to be the global optimum [41].

### Results

The final results of the simulation are summarized in Figure 8 and Table 5. The algorithm terminated after allocating the



**Figure 8:** On the upper left, the additional revenues  $\Pi_{spot,i}$  per month are plotted as a function of the spot instance power  $C_{spot,i}$  and as a function of the spot instance data rate  $R_{spot,i}$  on the upper right. The lower left shows the same demand elasticity as Figure 7. The circles indicate the optimized solutions for  $C_{spot,i}^*$ ,  $R_{spot,i}^*$  and  $p_{spot,i}^*$ .

complete available capacity. The upper left part of Figure 8 shows the spot revenues per month  $\Pi_{spot,i}(C_i)$  in thousands of USD (\$k) as a function of the capacity, whereas the upper right part shows the spot revenues as a function of the data rate  $\Pi_{spot,i}(R_{spot,i})$ . The lower left subplot is the demand elasticity from Figure 7. The location of the circles in all three subplots denote the optimal solutions obtained for  $C_{spot,i}^*$ ,  $R_{spot,i}^*$ , and  $p_{spot,i}^*$ . For the initialization of  $C_{spot,i} = 0$  the circles in the lower left plot match the allotment price points.

We can see in the upper left plot that the gradients  $\partial\Pi_i/\partial C_i$  are all equal at the optimum solution points. Due to the strong non-linearity of the link budget (see Eq. (5)), the functions and the gradients with respect to  $R$  are of a very different form. For example, trunking has a steep gradient in the beginning for power, whereas it is almost linear with the data rate. The linearity can be explained by the flattening out of the demand elasticity for a cumulative demand above 300 Mbps. According to the Shannon limit, this linearity becomes the non-linear behavior seen in the upper left plot. An optimization on the gradients of  $\partial\Pi/\partial R$  would lead to a different non-optimal solution. This observation underscores the uniqueness of a potential satcom RM system.

**Table 5: Results after optimization**

	$C_{spot,i}^*$ [W]	$R_{spot,i}^*$ [Mbps]	$p_{spot,i}^*$ [\$]	$p_{CIR,i} - p_{spot,i}^*$ %	$\Pi_{CIR,i}$ [\$k]	$\Pi_{spot,i}$ [\$k]	$\frac{\Pi_{spot,i}}{\Pi_{spot}}$
<b>Aviation</b>	0.14	39	202	33%	21	7.92	17%
<b>Backhauling</b>	0.13	92	80	20%	15	7.43	16%
<b>Maritime</b>	0.13	64	148	26%	20	9.47	21%
<b>Trunking</b>	0.17	234	70	30%	25	16.31	36%
<b>Vsat</b>	0.06	32	148	26%	10	4.75	10%
				$\Sigma_i$	<b>91</b>	<b>46</b>	
				$\Pi_{tot}$	<b>137</b>	<b>+34%</b>	

Overall, under the assumption that all customers buy the *spot instances* for the posted price and quantity, they add \$46k in revenues per month to the baseline of \$91k (see Table 5). This equals an increase of 34%. The greatest revenues (around one third) come from trunking, with vsat contributing the least at 10% (see column  $\Pi_{spot,i}/\Pi_{spot}$ ). The capacity is split almost equally between aviation, backhauling, maritime, and trunking (Vsat receives around half of that). The prices  $p_{spot,i}^*$  are between 20% and 33% lower than  $p_{CIR,i}$ . The additional data rate is the highest for trunking with 234 Mbps.

After optimization, these determined prices would be passed on to the booking system (see Figure 3). The price-based RM system would then offer the *spot instance* product to each customer segments, e.g. through a web portal.

### Discussion

Our objective of this Section was to illustrate the working principles of a possible realization of an end-to-end satcom

RM system given the framework developed in the previous Sections. We focused on three areas: available capacity estimator, model, and SLA contract add-on menu. We described a statistical capacity estimator, which makes use of customer multiplexing. The model is based on the link budget and we illustrate its non-linear behavior and the implications. For the add-on menu, we discussed one possible product: *spot instances* that are sold to the same customer for a discounted price, but the operator can quit service anytime. Under the assumption that the estimations about price elasticity are accurate and customers buy the *spot instances*, the operator's revenues can be boosted by 34%.

The biggest assumptions we made were around the add-on menu. The specification of the product has implications on the demand elasticities and the available capacity considered. While we assumed here that the elasticities are anchored around the price point of the long-term allotments, this is not necessarily the case or even known. An analogy are new flights for airlines. The demand curves are initially unknown and are learned over time.

Moreover, we considered the conservative scenario in which customers make full use of the purchased data rate of the spot instances. It is likely that they show similar behavior to the long-term allotments and leave considerable amount unused during the night. One approach can be to design a product that specifically makes use of this night drop (e.g. a discounted over-night backup plan for non-time critical data). Another approach is to store power during off-times and use during peak hours (depending on the technical capabilities of the satellite).

## 7. CONCLUSION

To summarize, in this paper we discussed the potential of Revenue Management (RM) for satcom operators. We motivated that several technological changes are underway, in particular digital payloads and phased arrays, that will change how operators control their satellites. Competition is likely to increase in the coming years and being able to manage demand and capacity effectively could become a competitive differentiator.

We showed that the satcom industry has promising characteristics for the implementation of a RM system: inflexible capacity, low marginal sales cost, perishable inventory, heterogenous customers, and variable and uncertain demand. We divide the satcom products into existing long-term allotments and future short-term add-ons. With the development of a Dynamic Resource Management (DRM) tool, operators will create the underlying information and data required for a potential RM system, hence reduce development risks and costs.

Drawing from analogies to the passenger airline and air cargo industries, we propose a price-based RM framework for satcom that has three unique features: available capacity forecaster, model, and SLA add-on menu.

Finally, we provided a proof-of-concept of the developed price-based RM framework. In the realization that we simulated, we introduced a spot instance product that makes use of half of the statistically available capacity and can significantly increase revenues compared to a long-term allotment only case.

We make the following contributions:

- Identifying RM as highly applicable for satcom operators to manage demand and capacity integratedly
- Proposing the first RM framework for satcom operators with the three unique features: available capacity forecaster, model, and SLA add-on menu
- Proofing through simulation that our RM framework can deliver a significant increase in revenues

Since this is the first paper on RM for satcom operators, it opens many unanswered questions that require further research and deeper analysis. We hope that the framework of this paper is helpful for some of that future work.

Possible future research directions can be:

- Development of an analytical model for NGSO satellites
- Development of a more exhaustive parameterized SLA add-on menu
- Modelling of the implication of storing power on the available capacity estimator
- Conducting a pilot study for validation of the framework

#### ACKNOWLEDGMENT

This work was supported by SES. The authors want to thank SES for their input to this paper and their financial support.

#### REFERENCES

- [1] T. Canada, "Telesat Ka-band NGSO constellation FCC filing SAT-PDR-20161115-00108," 2018. [Online]. Available: [http://licensing.fcc.gov/myibfs/forwardtopublictabaction.do?file\\_number=SATPDR2016111500108](http://licensing.fcc.gov/myibfs/forwardtopublictabaction.do?file_number=SATPDR2016111500108)
- [2] W. S. Limited, "OneWeb Ka-band NGSO constellation FCC filing SAT-LOI-20160428-00041," 2018. [Online]. Available: [http://licensing.fcc.gov/myibfs/forwardtopublictabaction.do?file\\_number=SATLOI2016042800041](http://licensing.fcc.gov/myibfs/forwardtopublictabaction.do?file_number=SATLOI2016042800041)
- [3] S. E. Holdings, "LLC, SpaceX Ka-band NGSO constellation FCC filing SAT-LOA-20161115-00118," 2018. [Online]. Available: [http://licensing.fcc.gov/myibfs/forwardtopublictabaction.do?file\\_number=SATLOA2016111500118](http://licensing.fcc.gov/myibfs/forwardtopublictabaction.do?file_number=SATLOA2016111500118)
- [4] Viasat. "Going Global - Viasat-2 and the Viasat-3 Platform Will Take Our Service Around the World." <https://www.viasat.com/news/going-global> (accessed October, 2018).
- [5] SES. "Exponentially More Opportunities With O3b mPOWER." <https://www.ses.com/networks/o3b-mpower> (accessed October, 2018).
- [6] SES. "SES Selects Arianespace for Launch of SES-17." <https://www.ses.com/press-release/ses-selects-arianespace-launch-ses-17> (accessed October, 2018).
- [7] M. Guerster, J. J. G. Luis, E. Crawley, and B. Cameron, "Problem representation of dynamic resource allocation for flexible high throughput satellites," presented at the IEEE Aerospace, Big Sky, MT, 2019.
- [8] A. I. Aravanis, B. S. MR, P.-D. Arapoglou, G. Danoy, P. G. Cottis, and B. Ottersten, "Power allocation in multibeam satellite systems: A two-stage multi-objective optimization," *IEEE Transactions on Wireless Communications*, vol. 14, no. 6, pp. 3171-3182, 2015.
- [9] A. Paris, I. d. Portillo, B. Cameron, and E. Crawley, "A Genetic Algorithm for Joint Power and Bandwidth Allocation in Multibeam Satellite Systems," presented at the IEEE Aerospace, Big Sky, Montana, 2019.
- [10] Y. He, Y. Jia, and X. Zhong, "A traffic-awareness dynamic resource allocation scheme based on multi-objective optimization in multi-beam mobile satellite communication systems," *International Journal of Distributed Sensor Networks*, vol. 13, no. 8, 2017.
- [11] N. Pachler, J. J. G. Luis, M. Guerster, E. Crawley, and B. Cameron, "Allocating Power and Bandwidth in Multibeam Satellite Systems using Particle Swarm Optimization," presented at the 2020 IEEE Aerospace Conference, Big Sky, Montana, 2020.
- [12] F. R. Durand and T. Abrão, "Power allocation in multibeam satellites based on particle swarm optimization," *AEU-International Journal of Electronics and Communications*, vol. 78, pp. 124-133, 2017.
- [13] H. Wang, A. Liu, X. Pan, and J. Yang, "Optimization of power allocation for multiusers in multi-spot-beam satellite communication systems," *Mathematical Problems in Engineering*, vol. 2014, 2014.
- [14] J. J. G. Luis, M. Guerster, I. del Portillo, E. Crawley, and B. Cameron, "Deep Reinforcement Learning Architecture for Continuous Power Allocation in High Throughput Satellites," *arXiv preprint arXiv:1906.00571*, 2019.
- [15] G. Cocco, T. De Cola, M. Angelone, Z. Katona, and S. Erl, "Radio Resource Management Optimization of Flexible Satellite Payloads for DVB-S2 Systems," *IEEE Transactions on Broadcasting*, vol. 64, no. 2, pp. 266-280, 2018.
- [16] X. Hu, S. Liu, R. Chen, W. Wang, and C. Wang, "A Deep Reinforcement Learning-Based Framework for Dynamic Resource Allocation in Multibeam Satellite Systems," *IEEE Communications Letters*, vol. 22, no. 8, pp. 1612-1615, 2018.
- [17] S. Liu, X. Hu, and W. Wang, "Deep Reinforcement Learning Based Dynamic Channel Allocation Algorithm in Multibeam Satellite Systems," *IEEE ACCESS*, vol. 6, pp. 15733-15742, 2018.
- [18] U. Park, H. W. Kim, D. S. Oh, and B. J. Ku, "A dynamic bandwidth allocation scheme for a multi-spot-beam satellite system," *Etri Journal*, vol. 34, no. 4, pp. 613-616, 2012.
- [19] J. Camino, C. Artigues, L. Houssin, and S. Mourgues, "Mixed-integer linear programming for multibeam satellite systems design: Application to the beam layout optimization," in *2016 Annual IEEE Systems Conference (SysCon)*, 18-21 April 2016 2016, pp. 1-6, doi: 10.1109/SYSCON.2016.7490613.
- [20] A. Kyrgiazos, B. Evans, and P. Thompson, "Irregular beam sizes and non-uniform bandwidth allocation in HTS multibeam satellite systems," in *31st AIAA International*

*Communications Satellite Systems Conference (ICSSC)*, 2013.

- [21] N. Pachler, M. Guerster, I. del Portillo, E. Crawley, and B. Cameron, "Static beam placement and frequency plan algorithms for LEO constellations," *International Journal of Satellite Communications and Networking*, 2020.
- [22] S. E. Kimes, "The basics of yield management," *Cornell Hotel and Restaurant Administration Quarterly*, vol. 30, no. 3, pp. 14-19, 1989.
- [23] K. T. Talluri and G. J. Van Ryzin, *The theory and practice of revenue management*. Springer Science & Business Media, 2006.
- [24] B. C. Smith, J. F. Leimkuhler, and R. M. Darrow, "Yield management at American airlines," *Interfaces*, vol. 22, no. 1, pp. 8-31, 1992.
- [25] P. P. Belobaba, "Survey Paper—Airline yield management an overview of seat inventory control," *Transportation science*, vol. 21, no. 2, pp. 63-73, 1987.
- [26] P. Belobaba, A. Odoni, and C. Barnhart, *The global airline industry*. John Wiley & Sons, 2015.
- [27] P. P. Belobaba, "OR practice—application of a probabilistic decision model to airline seat inventory control," *Operations Research*, vol. 37, no. 2, pp. 183-197, 1989.
- [28] W.-C. Chiang, J. C. Chen, and X. Xu, "An overview of research on revenue management: current issues and future research," *International journal of revenue management*, vol. 1, no. 1, pp. 97-128, 2007.
- [29] L. R. Weatherford and S. E. Bodily, "A taxonomy and research overview of perishable-asset revenue management: Yield management, overbooking, and pricing," *Operations research*, vol. 40, no. 5, pp. 831-844, 1992.
- [30] SES, "Personal communication," ed, 2019.
- [31] NSR, "Flat Panel Satellite Antennas," 2019, vol. 4th Edition. [Online]. Available: <https://www.nsr.com/research/flat-panel-satellite-antennas-4th-edition/>
- [32] M. Wedel and W. A. Kamakura, *Market segmentation: Conceptual and methodological foundations*. Springer Science & Business Media, 2012.
- [33] SES. *Proprietary data*.
- [34] J. S. Billings, A. G. Diener, and B. B. Yuen, "Cargo revenue optimisation," *Journal of Revenue and Pricing Management*, journal article vol. 2, no. 1, pp. 69-79, April 01 2003, doi: 10.1057/palgrave.rpm.5170050.
- [35] R. G. Kasilingam, "Air cargo revenue management: Characteristics and complexities," *European Journal of Operational Research*, vol. 96, no. 1, pp. 36-44, 1997.
- [36] B. Becker and N. Dill, "Managing the complexity of air cargo revenue management," *Journal of Revenue and Pricing Management*, journal article vol. 6, no. 3, pp. 175-187, September 01 2007, doi: 10.1057/palgrave.rpm.5160084.
- [37] B. Slager and L. Kapteijns, "Implementation of cargo revenue management at KLM," *Journal of Revenue and Pricing Management*, vol. 3, no. 1, pp. 80-90, 2004.
- [38] P. Belobaba, A. Odoni, and C. Barnhart, *The Global Airline Industry*. 2009.
- [39] G. Maral and M. Bousquet, *Satellite communications systems: systems, techniques and technology*. John Wiley & Sons, 2011.
- [40] Amazon. "EC2 Spot Instances." <https://aws.amazon.com/ec2/spot/> (accessed October, 12, 2019).
- [41] A. Federgruen and H. Groenevelt, "The greedy procedure for resource allocation problems: Necessary and sufficient conditions for optimality," *Operations research*, vol. 34, no. 6, pp. 909-918, 1986.

## BIOGRAPHY



**Markus Guerster** is a 3<sup>rd</sup> year Ph.D. student in the System Architecture Lab/Engineering Systems Lab of the Department of Aeronautics and Astronautics at MIT. He works on dynamic resource allocation for communication satellites with particular focus on Revenue Management for operators. He received a M.Sc. degree (2017) with high distinction in Aerospace Engineering as well as a B.Sc. degree (2015) in Mechanical Engineering from the Technical University of Munich.



**Joël Grotz** graduated from the University of Karlsruhe and the Grenoble Institute of Technology in electrical engineering and holds a Ph.D. degree in Telecommunications from the Royal Institute of Technology (KTH), Stockholm. He worked at the Technology Labs Department at Newtec Cy in Belgium, leading development projects and working on modems and novel satellite transmission techniques from 2012 to 2015. Currently he is with SES Engineering working on future satellite communication projects.



**Peter P. Belobaba** is Principal Research Scientist in the International Center for Air Transportation at MIT, where he teaches graduate subjects on The Airline Industry, Airline Management, and Air Transportation Operations Research. He is Program Manager of MIT's Global Airline Industry Program and Director of the PODS Revenue Management Research Consortium. Dr. Belobaba holds an MS in Transportation and a Ph.D. in Flight Transportation Systems from MIT. He is a lead author and editor of the recently released book, *The Global Airline Industry, 2nd Edition*. He has worked as a consultant on the development, testing and implementation of pricing, revenue management and distribution systems at over fifty airlines and other companies worldwide. In 2016, Dr. Belobaba was awarded the INFORMS Impact Prize for his "pivotal role in the creation and wide-spread adoption of revenue management".



**Edward F. Crawley** is the Ford Professor of Engineering, and a Professor of Aeronautics and Astronautics at MIT. He has served as the founding President of the Skolkovo Institute of Science and Technology (Skoltech) in Moscow, the founding Director of the MIT Gordon Engineering Leadership Program, the Director of the Cambridge (UK) MIT Institute and the

*Head of the Department of Aeronautics and Astronautics at MIT. Dr. Crawley is a Fellow of the AIAA, the Royal Aeronautical Society (UK) and a member of the International Academy of Astronautics. He is a member of five national academies: in Sweden, the UK, China, Russia and the US. He received an S.B. (1976) and an S.M. (1978) in aeronautics and astronautics and a Sc.D. (1981) in aerospace structures, all from MIT, and has been awarded two degrees of Doctor Honoris Causa. Crawley's research has focused on the architecture, design, and decision support and optimization in complex technical systems subject to economic and stakeholder constraints. His recent book – System Architecture: Strategy and Product Development for Complex Systems – was published by Pearson (2016).*

---



**Dr. Bruce Cameron** is the Director of the System Architecture Group at MIT and a consultant on platform strategies. At MIT, Dr. Cameron ran the MIT Commonality study, a 16 firm investigation of platforming returns. Dr. Cameron's current clients include Fortune 500 firms in high tech, aerospace, transportation, and consumer goods. Prior to MIT, Bruce worked as an engagement manager at a management consultancy and as a system engineer at MDA Space Systems, and has built hardware on orbit. Dr. Cameron received his undergraduate degree from the University of Toronto, and graduate degrees from MIT.