So Many Beams, So Little Time: Revenue Management in the Next Generation of Flexible Communication Satellites

Markus Guerster[a], Edward Crawley[b], Sergi Aliaga[b], Bruce Cameron[b]

[a]*SES*
1129 20th St NW STE 1000
Washington, DC 20036
Markus.Guerster@ses.com
[b]*Massasachusetts Institute of Technology*
77 Massachusetts Av. 33-409
Cambridge, MA 02139
{*crawley, aliaga, bcameron*}*@mit.edu*

Abstract

In this paper we present a proof of concept of how satellite operators can monetize the available resources granted by the latest advancements in satellite payloads, which provide customers with unprecedented levels of flexibility when it comes to resource allocation. The degrees of flexibility range from adjustable power only to full control over frequency assignment, routing, beam pointing and beam shape.
Furthermore, the current satcom landscape is expected to be governed by highly variable demand, making old and new satellite operators face a new challenge in how to manage demand and capacity. Previous literature has pointed out the similarities of the current satcom market with highly established industries, like Airline or Car Rental, remarking its huge potential for a new Revenue Management strategy.
We present two different strategies that can leverage a satcom-adapted Revenue Management framework to lift the revenues of satcom operators by more than 20%. Both strategies are based on monetizing the available capacity through existing customers: either selling more capacity or selling additional products. We outline the complete process: we first detail the general and specific assumptions of each monetizing approach and the pricing optimization strategy employed. We then set up a simulation that contains actual demand data and quantify the benefit of dynamic power allocation in terms of available capacity, compared to fixed-by-design allocation. Finally, we present the results of applying the two mentioned strategies and provide a comparison between them.

Keywords: Revenue Management, Satellite Communications, Flexible payloads, Resource Allocation

Contents

1. Introduction

The flexibility of satellites have been increasing for the last 20 years. Transparent digital processors and phased antenna arrays enable this flexibility and allow for the dynamic allocation of resources. The degree of flexibility ranges from adjustable power only to full control over frequency assignment, routing, beam pointing, and beam shape. This flexibility motivates a need for dynamic ways to adapt to variable demand and free-up resources whenever possible.

Since the first high-throughput satellite (HTS) Thaicom 4 [1] (throughput of 45Gbps), HTS have been continuously increasing its combined throughput – mostly thanks to the introduction of spot beams. One example is SES's latest O3b mPOWER constellation [2]. It consists of up to 11 Middle Earth Orbit (MEO) satellites providing a combined throughput of over 1 Tbps [3] – almost two orders of magnitude more than Thaicom 4. Such unprecedented throughput is achieved by means of over 4000 spot beams which are dynamically controllable by digital communication payloads and multi-beam phased array antennas [4, 5].

There is extensive research on how to exploit these advantages in flexibilities for an optimal and dynamic resource allocation. Most of the work in that area falls in the category of power allocation. Aravanis [6], Paris [7], He [8], Pachler [9], and Durand [10] use metaheuristics. Convex optimization is applied by Wang [11] and Deep Reinforcement Learning by Garau [12]. The frequency assignment problem is less studied with Paris [7] and Cocco [13] applying metaheuristics, Hu and Liu [14, 15] and Garau [16] Deep Reinforcement Learning, and Park [17] a binary search heuristic. Beam pointing and beam shape are addressed by Camino [18] with linear programming and by Kyrgiazos [19] with a self-organizing map. Pachler et al. [20] developed two heuristic algorithm to solve the frequency assignment and beam pointing problem with particular focus on LEO constellations. Some of these techniques are even implemented for autonomous satellite onboard mission activity planning ([21], [22])

Despite all those efforts in freeing up resources, little research is addressing how satellite operators could monetize them. By reviewing previous literature on demand growth forecasts for satellite communications, in the following section we identify the need of a study on how satellite communication (satcom) operators can monetize those available resources provided by modern flexible payloads and state-of-the art Dynamic Resource Management techniques.

2. Literature review

Traditional economic studies in satcom tend to address the development stages and economic impact of satellite technologies on its upstream value-chain ([23–27]) but close to none talk about the economics of the assets once they are out in space.

Numerous reports predict a high growth for bidirectional and bursty demand, which thanks to the new advancements presented before, becomes a highly appealing market for satellite operators. In 2017, Morgan Stanley released an investment report predicting the future of the space economy until 2040 [26]. They expect the global demand for data to grow at an exponential rate due to an increase in the global population, autonomous cars, Internet of things, artificial intelligence, virtual reality, and video applications. Several other sources [28–32] expect similarly high rates of growth with an often-cited doubling of data volume every two years. Morgan Stanley [26] and a 2019 Northern Sky Research (NSR) report [27] believe that the satellite broadband market can leverage that growth. Morgan Stanley forecasts an increase in broadband revenue opportunity of two orders of magnitude by 2040. The NSR sees a tripling of the number of satcom consumer broadband subscribers until 2027 for a low growth scenario and a tenfold increase for the high growth scenario. An SES report [33] cites a particularly high growth for the mobility sector, in particular for airplanes and maritime, which aligns with the new market opportunities considered by Morgan Stanley: airplanes, maritime, trains, and trucks, and automobiles. Kota [34] and Farserotu [35] specifically discuss the mobility sector and the integration of satellite networks. As an additional future market, Hosseini [36] reviewed how satellites can be integrated into a 5G landscape to play an essential role in providing control and data links to Unmanned Aerial Vehicles (UAVs).

In conclusion, with the expected high growth of demand for data, the need for satellite communication is expected to grow as well [26, 27, 33]. That is especially true for the mobility sector, where terrestrial alternatives perform more poorly [26, 33, 36]. However, there is no public study or market dynamics model to understand the dynamics of the satcom market, nor work on Revenue Management (RM) for satcom.

The purpose of this paper is to rethink the Revenue Management framework in order to cover the four layers described by Talluri in [37]: data input, estimation and forecasting, optimization and control. Before presenting the

specifics of the modeling methodology employed in section 4, in the next section we describe the two proposed ways to monetize the available capacity provide by the resource management strategies mentioned before, as well as some general and specific assumptions for each one of them. Additionally, we present the general pricing optimization procedure to implement those two approaches.

3. Monetizing the available capacity through existing customers

The two revenue lift strategies that we consider are based on selling more capacity to existing customers, which means that their demand needs to increase. That can be achieved by selling more volume of the existing Service Level Agreements (SLAs) (Section 3.1, referred to as demand for use in terrestrial telecommunication [38]), or provide customers with additional products (Section 3.2). The other option would be to contract new customers through price adjustments, unlocking the affordability elasticity of the segments (referred to as access demand in terrestrial telecommunication [38]). However, we do not cover this possibility in this paper. We make the following general assumptions for the two strategies considered:

- Demand at the same price remains constant throughout the simulation, with no growth nor shrinking.
- Price does not affect the entrance of new users (no affordability elasticity)
- Customers do not switch operator when prices are changed (no competition)
- Elasticity is known per customer
- Price is the only lever to control demand.
- Separate prices can be charged per customer.
- Recomputing the frequency plan does not improve the result and power adjustments are sufficient for the changes in demand

The specific assumptions for each of the two approaches are described in the following subsections.

3.1. Selling more capacity through existing SLAs: Specific approach and assumptions

The vast majority of current satcom operator's customers are contracted by the Classical SLA with a price for the Committed Information Rate (CIR). Hence, we assume that all customers have this contract type, with the baseline being their current price points for the CIR. Selling more quantity of the Classical SLA translates into contracting a higher CIR. We make the assumptions here that the usage behavior of users changes slower than users' SLA are re-negotiated to allow for safe over-subscriptions. If an operator has 80 users with an average SLA duration of 2 years, then on average, a little over one SLA expires per week. If the total traffic of the users does not increase more than the SLA's CIR per week, then the operator can safely oversubscribe.

We model the dependency between the CIR and the price with a price elasticity for the demand-for-use. Since most terminals have their own CIRs, we compute the elasticity curve on a user terminal level. We build a demand elasticity regression by means of a power/log-linear curve with constant elasticity:

$$R = a \cdot p^b \text{ or } p = \left(\frac{R}{a}\right)^{\frac{1}{b}} \tag{1}$$

where a and b are fitting parameters with b having the specific meaning of the elasticity. R and p are the data rate and price respectively. When applying the logarithmic function to the equation, we get the linear Eq. (2) for which we can compute the coefficients by minimizing the residual sum of squares with linear approximation.

$$\log(R) = \log(a) + b \cdot \log(p) \tag{2}$$

We use the log-linear relationship from Eq.(1), where the current price point p_{CIR} for the CIR determines the parameter a through the following equation:

$$a = CIR/p_{CIR}^b \tag{3}$$

Figure 1 illustrates three demand elasticity (-2, -1, -0.5) for an example customer with a CIR of 100~Mbps for a price of $200/month/Mbps. For the unit elasticity -1, reducing prices to $100 increases the demanded data rate to 200~Mbps (and therefore no change in revenues). If the demand is more elastic with -2, the same price leads to twice the demanded data rate, 400 Mbps. The opposite holds for an inelastic demand of -0.5; the CIR goes down to 141.5~Mbps.

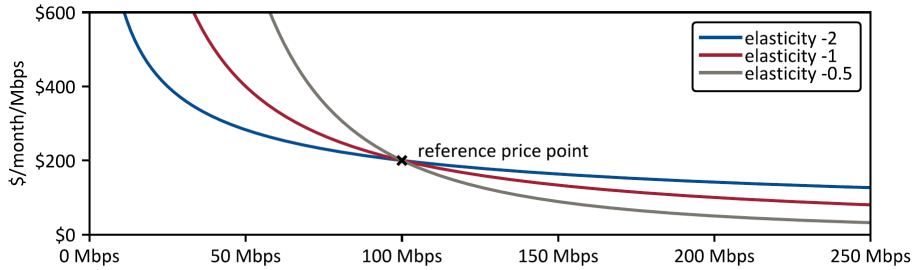

Figure 1: Example of a -2, -1, and -0.5 price elasticity for a customer with a 100 Mbps and $200/month/Mbps

For the analysis, we compare three different elasticity scenarios. For two of them, all customers have the same elasticity that is either elastic with -2, or unit elastic with -1. In the third case, the elasticity is different for each customer varying between -2 and -1. The assignment of elasticity to a customer is random from this interval.

3.2. Selling additional products: Specific approach and assumptions

In contrast to the previous approach, this Section discusses the selling of the available capacity through additional products. These products can be novel SLAs. One example here is the spot instance that an operator might want to use an add-on. It gives the customer additional flexibility to purchase capacity, which is especially valuable for customers who cannot predict their traffic well in advance. We leave the existing SLAs untouched in this approach, and therefore, assume that the products are priced and differentiated accordingly such that no internal cannibalization occurs.

One central question is how operators can price these new products and what the demand elasticity is. Since almost no information is available, we model the price point and elasticity explicitly with three parameters (two for the price point and one for the elasticity) and conduct a broad sensitivity analysis.

The first parameter that we introduce is the fraction of revenues $k_{rev,fract}$ that the additional product delivers compared to the existing Classical SLA. The second parameter is the adjustment parameter k_{adj} that shifts the reference point through which the elasticity curve is defined (see Figure 2). The following three Eqs. (4) - (6) describe the mathematical relationships between these two parameters, and the existing SLA attributes p_{CIR} and CIR. The result is the reference point defined by $p_{ref,add}$ and $R_{ref,add}$. We take the square root of $k_{rev,fract}$ in the first two equations, so that $k_{rev,fract}$ affects the additional revenues in Eq. (6) linearly.

$$p_{ref,add} = p_{CIR} \cdot k_{adj} \cdot \sqrt{k_{rev,fract}} \tag{4}$$

$$R_{ref,add} = CIR \cdot \frac{1}{k_{adj}} \cdot \sqrt{k_{rev,fract}} \tag{5}$$

$$\Pi_{add} = p_{ref,add} \cdot R_{ref,add} = k_{rev,fract} \cdot p_{CIR} \cdot CIR \tag{6}$$

To capture the range of possible prices for the new products, we set up nine combinations of the two parameters, $k_{rev,fract}$ and k_{adj} (see Table 1).

The fractions of revenues ranges in three steps: 10%, 20%, and 50%. The second parameter k_{adj} has the steps: 0.2, 0.5, and 1.

Figure 2 provides a visual understanding of the impact of different values, including the elasticities. The black dashed line is the unit elasticity from the previous subsection example in Figure 1 with its reference price point of $200/Mbps/month for 100 Mbps. We set $k_{rev,fract}$ to 20% and vary the elasticity from -0.5 to -2.

	Comb. 1	Comb. 2	Comb. 3	Comb. 4	Comb. 5	Comb. 6	Comb. 7	Comb. 8	Comb. 9
$k_{rev,fract}$	0.1	0.1	0.1	0.2	0.2	0.2	0.5	0.5	0.5
k_{adj}	0.2	0.5	1	0.2	0.5	1	0.2	0.5	1

Table 1: Test plan for testing of the sensitivity to the reference price point of the additional product. Each combination is simulated for four elasticity cases: -2, random between -2 and -1, -1, and -0.5.

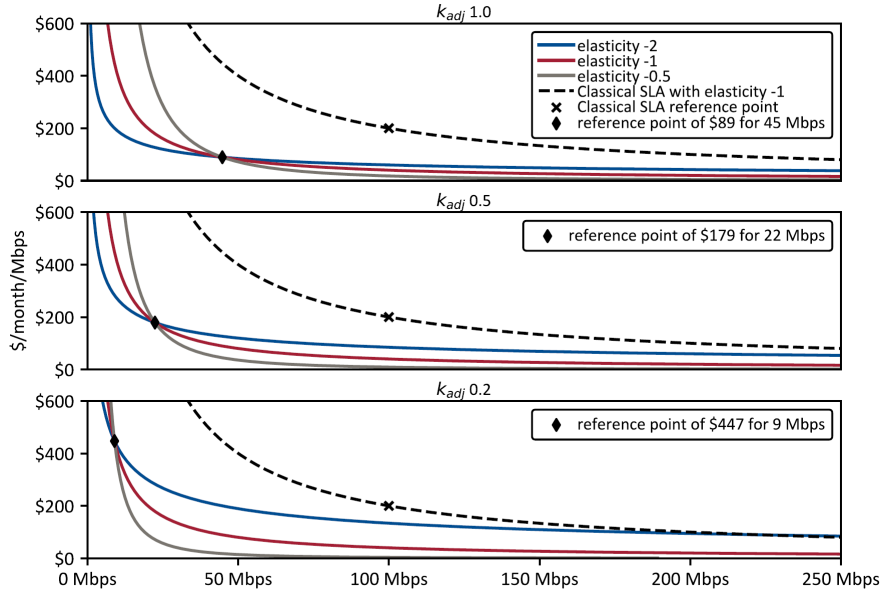

Figure 2: Behavior of the price elasticity function as a function of the reference point and elasticity parameter. The plots are shown for a price 20% that of the Classical SLA for the unit elasticity.

By definition, the red unit elasticity lines are independent of the parameter k_{adj}. That can be seen through Eq. (3). For $b = -1$, the equation for a becomes $a = CIR \cdot p_{CIR}$, and since k_{adj} is in the numerator for p_{add} and in the denominator for R_{add} (see Eqs. (4) . (5)), k_{adj} cancels out. However, elasticities different from the unit elasticity are dependent on k_{adj}. A higher value implies that the inelastic effect occurs at higher Mbps, whereas the elastic effect appears at a lower price (can be seen at the behavior of the blue line for large Mbps numbers). This behavior might be a good approximation for a product that a larger quantity is sold for a lower price. In contrast, for a lower value of k_{adj}, the inelastic case constrains the Mbps strongly while in the elastic case, the operator can sell larger quantities for a higher price. An example product has an initial high price for a smaller quantity.

3.3. General pricing optimization approach

As presented at the beginning of this section, we now describes the elements of the Revenue Management optimization that are common across the two analyses of this article.

The Demand and Resource management parts of the satcom RM framework need to exchange information (see Figure 3). The pricing optimization computes new prices that updates the expected traffic. The resource management allocates resources for this new traffic and returns the used and available capacity to the demand management part. Hence, the pricing algorithm calls the resource allocation simulation every iteration.

Figure 3 depicts the process that is common for the two analyses. We initialize the optimization with the baseline allocation from Section 5. The iteration starts with an estimation of the customer and market price elasticities.

The results inform the pricing policy optimization, which computes its next iteration of prices. Based on these, the optimization updates the expected traffic volume and builds customer usage history. Depending on the significance of the change in the traffic, the resource allocation adjusts the power levels or additionally revises the frequency plan.

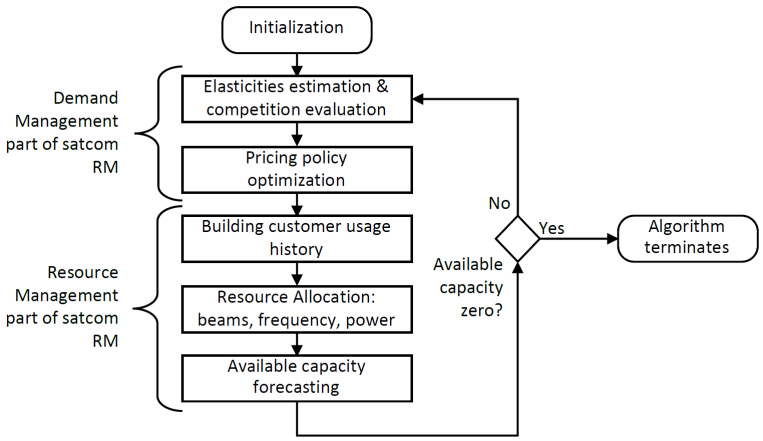

Figure 3: General approach for integrated optimization of pricing and resource allocation.

Based on the results, the available capacity forecaster determines the used capacity and subtracts it from the total capacity. If the available capacity is zero (or as close as possible given the discretization), the algorithm terminates. Otherwise, the next iteration begins with the elasticities' estimation and evaluation of the competition.

3.3.1. Sorted list of prices and binary search

Due to the monotonic behavior of the price elasticities and the link budget, lowering prices does not decrease demand, and increasing prices does not increase demand. With this characteristic, we can build a sorted list, which is, when sorted by prices, also inversely sorted by demand, and therefore by used capacity (lower price yields more demand, more used capacity, less available capacity). Table 2 illustrates a schematic example for a satellite with a single user i. The maximum capacity is 6 Watts, bandwidth 50 MHz, and C_{used} is computed proportionally with the Shannon limit according to $C_{used} = 2^{R_i/50MHz}$. The line with bold font indicates the price for which the available capacity is minimized (and revenues maximized if demand is elastic).

p_i [$/Month/Mbps]	R_i [Mbps]	C_{used} [W]	C_{avail} [W]
283	50	2.00	4.00
258	60	2.30	3.70
239	70	2.64	3.36
224	80	3.03	2.97
211	90	3.48	2.52
200	100	4.00	2.00
191	110	4.59	1.41
183	**120**	**5.28**	**0.72**
175	130	6.06	-0.06
169	140	6.96	-0.96
163	150	8.00	-2.00

Table 2: Schematic example of the ordered list of price p_i and the resulting demand R_i for a single user i (elasticity is -2 at $p_i = 200$ and $R_i = 100$). Used capacity computed through the Shannon limit and available capacity with a maximum capacity C_{max} of 6 Watts.

Building the list of declining prices is the outcome of the pricing policy optimization and computationally orders of magnitude cheaper than the mapping between prices and used capacity. Since the list is sorted, we use a binary search to find the element where the available capacity is zero, i.e., the bold line in Table 2. An element comprises a set of prices for each customer. The logarithmic complexity with the length of the list N_l is $O(log(N_l))$. Hence the algorithm only makes a few evaluations of the available capacity for a longer array of prices (e.g., three comparisons for $n = 1,000$, and four for $n = 10,000$). A more detailed explanation and the pseudo-code of the algorithm can be found in Appendix A

3.3.2. Algorithmic approach for computing the ordered list of prices

For the computation of the ordered list of prices, we compare mainly two algorithms. The first algorithm is an *equally decreasing heuristic*, which might be considered as closest to the current manual approach. When capacity is available, the operator discounts the CIR prices for all customers by the same percentage points Δp. The second algorithm is a *gradient optimizer based on marginal revenues*. It considers the complete chain from price elasticities to link budgets.

4. Methodology

In this section we describe the framework and its different elements used to obtain the results presented in Section 6. We start with a short description of the proposed Revenue Management framework, followed by a description of the input data and model elements that are the same to all simulations of this paper. That includes details on user terminals, gateways, constellation, traffic usage, and elasticities. All of them are based on an O3b mPOWER like MEO constellation.

4.1. Proposed Revenue Management Satcom Architecture

The Revenue Management framework proposed is built up on the one presented in (markus paper) and is depicted in Figure 4. It covers the four layers described by Talluri in [37] (data input, estimation and forecasting, optimization and control). However, this particular framework makes a distinction between stages purely dedicated to Demand Management, which are common to other RM frameworks (airlines, rental car, hotel,...), and stages dedicated to the Resource Management, which are a unique extension for the satcom RM framework and mostly cover the resource allocation processes mentioned before.

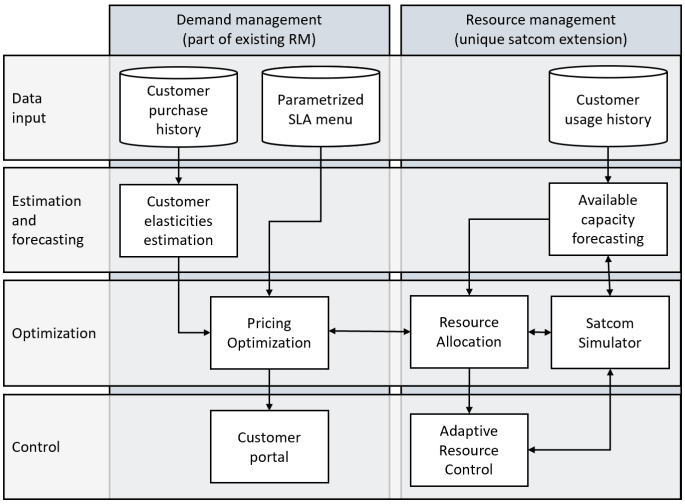

Figure 4: Overview of our proposed satcom RM framework

4.2. User traffic data

The aim of this subsection is to provide a description of the available data from selected SES's current users (note that we use in this article user and customer interchangeably). For each customer, we have the following features:

- Segment (encoded)
- CIR
- Monthly recurring revenues from which we compute the price in $/Mbps/Month
- SLA status: active, ended, closed but not yet active

- SLA service start and end
- Location information: approximate latitude and longitude, country name, and iso-a2 code
- Terminal size from which we derive the G/T (see Table 3)
- Preprocessed usage data: mean usage over 24h in 1-minute granularity and uncertainty expressed in a sigma value for the average of the means

terminal size $[m]$	EIRP $[dB]$	G/T $[dBK]$	G $[db]$
0.3	46	9	32.7
0.6	52	15	38.7
1.2	61	21	44.7
2.4	70	27	50.7
4.5	80	35	58.7
7.3	85	40	63.7

Table 3: Terminal EIRP, G/T, and G assumptions based on the terminal diameter

Data is sanitized to comply with confidentiality agreements. In particular, we adjust all prices by multiplying with a constant number and provide approximative information of the latitude and longitude. We use this dataset to build our two databases, *customer usage history*, and *customer purchase history* (see Figure 4). We describe the usage history here below.

We consider 80 *active* worldwide distributed user terminals with an accumulated CIR of 22.7 Gbps. The mean CIR per user terminals is 296 Mbps, with the smallest having 6 Mbps and the largest 1500 Mbps. These user terminals belong to four different *customer groups*, encoded as A, B, C, and D. Figure 5 depicts a histogram of the CIR and the size of each segment based on their CIR. The current active customers of SES have most of the terminals below 100 Mbps, although some of them have higher CIRs, with considerable density of user terminals above 700 Mbps (see the left plot of Figure 5).

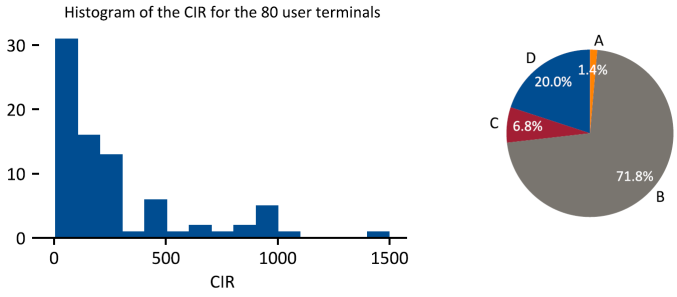

Figure 5: Statistical description of the 80 users

The right plot of Figure 5 shows the relative size of the four customer groups measured by the CIR. Group A is the smallest, with only accounting for 1.4% of the total CIR. Group B is the largest, with 72%, followed by group D and C with 20% and 7%, respectively.

4.3. Elasticity

From sections 3.1 and 3.2 we showed that results are sensitive to the elasticity. Therefore, understanding the elasticities in each segment is crucial. On a first approach, we estimated them based on past purchases. However, we deem the results from the regression as inaccurate, and we, therefore, did not use them for our analyses. As a second approach, we review the academic literature on elasticities in the telecommunication industry with the belief that the results are transferable to satcom (as there is no academic literature on broadband satcom elasticities). We combine the outcomes from both approaches to find a range of elasticities for sensitivity simulations in this Section.

From a customer perspective, the terrestrial telecommunication industry is the closest analogy to satcom. Hence, we review here selected econometric papers on the price elasticity estimation of customers in the telecommunication industry. Our goal is to find a typical range.

Aldebert et al. [39] use a translogarithmic indirect utility function to analyze residential demand for different traffic destinations (2004). Besides the price elasticity, they found that there is a considerable income elasticity. In analogy to satcom, this suggests that geographical segmentation is critical.

Comparing Aldebert et al. [39] results with other works, some results are different. For local traffic, the authors found an elasticity of -1.44 compared to -0.09 by Park et al. [40] and -0.88 by Wolak [41]. For long-distance traffic, Gatto et al. [42] compute -0.72, while Wolak [41] has the highest elasticity with -2.07. These numbers compare to -1.33 for national traffic from Aldebert [38]. They attribute the discrepancies to considering short- (under one year) vs. long-run (over one year) elasticities and different modeling approaches. For international traffic, Aldebert et al. estimate a fairly inelastic demand with -0.11, while Garin-Muñoz and Perez-Amaral [43] compute values between -1.31 and -0.32 varying between countries.

Ouwersloot and Rietveld [44] focus their review on the distance dependency of the price elasticity in telecommunication. They found that a doubling in the distance leads to a 0.07 increase in the elasticity. They categorized their review into local, national, and international calls with ranges from -0.05 to -0.75, -0.24 to -2.57, and -0.03 to -2.19, respectively. These scopes are similar to the ones discussed by Aldebert [38] and further illustrate the challenge associated with computing elasticities.

Hackl et al. [45] concentrate on the international telecommunication between Sweden and Germany, the United Kingdom, and the USA. They conclude that short and long-term price elasticities are in a similar range to other studies. The short-term elasticity ranges from -0.09 to -0.98, and the long-term from -0.19 to -1.61.

In summary, we reviewed multiple econometric papers on the price elasticity of demand in telecommunication and found that the difference in the reported numbers is considerable. With our regression on satcom elasticity, we found similar challenges. These observations motivate us to treat the elasticity as an uncertain parameter on which we conduct several sensitivity analyses. By combining the values from both approaches, we pick a range from -2 to -0.5 that covers elastic, unit elastic, and inelastic demand.

4.4. Gateways

We assume seven gateways distributed around the world, as shown in Table 4. They are in Hawaii, Texas, Azores, Sau Paulo, Cape Town, Colombo, and Auckland. All gateways serve out of country traffic and have a terminal size of 4.5m, giving them a G/T of 35 dB (see Table 3).

Unique ID	Lat [deg]	Lon [deg]	country code [iso-a2]	country name [-]	terminal size [m]	in/out country [-]
GW_Hawaii	19.9	-155.6	US	United States of America	4.5	out
GW_Texas	30.6	-96.3	US	United States of America	4.5	out
GW_Azores	37.7	-25.7	PT	Portugal	4.5	out
GW_SauPaulo	-23.6	-46.6	BR	Brazil	4.5	out
GW_CapeTown	-34.4	18.3	ZA	South Africa	4.5	out
GW_Colombo	6.9	79.8	LK	Sri Lanka	4.5	out
GW_Auckland	-36.5	174.6	CK	New Zealand	4.5	out

Table 4: Overview of the location and sizes of seven gateways considered

4.5. Space Segment

Since SES's O3b legacy MEO constellation serves these customers, we consider for this analysis a space segment is similar to the new O3b mPOWER MEO constellation. Table 5 lists the detailed parameters.

	Value
Number of satellites	7 for MEO
Orbit	Equatorial without eccentricity at 8075 km (MEO)
Repeating ground track	every 6 hours
Half-cone angle Θ_{3dB}	0.7 deg
Peak antenna gain $G_{Tx,max}$	35 dB
Total available bandwidth	2 GHz
Number of beamchannels	200 with 10 MHz
Number of reuses	4 including L/R polarization
Maximum number of beams	800

Table 5: Parameters of the MEO constellation, which is O3b mPower like

From the table, the maximum number of beams is 800, with the 10 MHz beamchannels. In the following sections, we analyze how much dynamic allocation reduces resource consumption, i.e., how much capacity becomes available (Section 5). With this metric, we measure the improvement from dynamic resource allocation in terms of power. In Section 6, we present the results obtained by following the two specific approaches presented in Section 3.

5. Resource allocation for the baseline

The goal of this Section is to establish a baseline. We use the pre-processed input data about the user usage, gateways, and the MEO constellation from Section 4. To that purpose we establish the following assumptions:

- We provide each user with its beam making the step of grouping the user terminals unnecessary.
- The routing is based on a balanced allocation (Gateways with similar load factor between them).
- The frequency assignment includes the return downlinks to the gateways assuming that all downlinks are in the same frequency band.
- We compare three different power allocation scenarios: (1) fixed by design, i.e., constant per beam, (2) stationary for CIR, (3) dynamic for actual usage.

We depict the resulting resource allocation solution in Figure 6. The upper part of the figures shows the users and gateways distributed around the world. The olive lines are active beams between satellites and gateways. The bottom plot shows the frequency assignment for the seven MEO satellites. It can be seen that satellites over the Pacific see less demand. This factor, together with the suboptimality of the algorithm results in only a 25% utilization of the frequency spectrum overall.

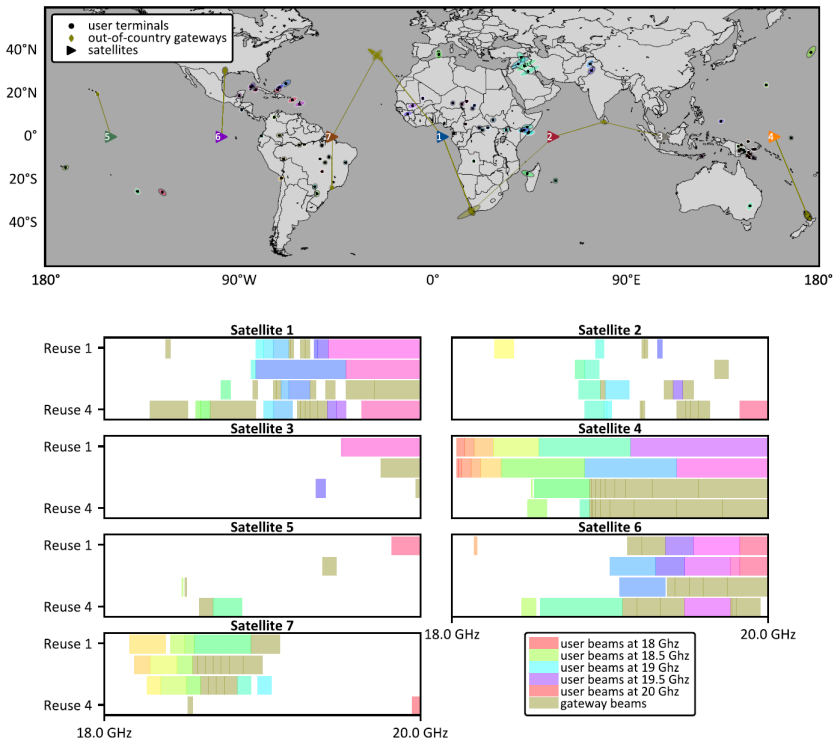

Figure 6: Resource Allocation solution after the first three steps of the resource allocation process (beam position, beam shape and frequency allocation) for the MEO O3b mPower like constellation. The usage of the frequency spectrum is on average 25%.

Using the solution of the first three steps of the resource allocation (beam position, beam shape and frequency allocation), we compute the power for three different scenarios:

1. **Fixed by design**. In this scenario, the satellite payload does not have any flexibility to adjust its power setting over time (traditional satcom). The design process determines the amount of transmitted power. We (optimistically) assume that the power is set to a constant level per beam meeting CIR at the worst case. The worst case is when the free space losses are the highest, which is when the user terminal and the satellite is separated the furthest. We assume in this scenario that beams can be activated and deactivated if not connect to a user.
2. **Stationary for CIR**. The second scenario considers a flexible payload where power is available as a pool. The satcom operator is conservative and provides the power levels that provided each user with their contracted CIR.
3. **Dynamic for actual usage**. In the last scenario, the satcom operator allocates dynamically the power that is needed by the user for any given time.

We simulate the power levels for all three scenarios with a 1-minute granularity over a 24-hour window and record the results. Figure 7 summarizes the aggregation of the data for satellite #3 (fixed by design in green, stationary for CIR in blue). The red line is the mean when power is allocated for the actual usage, and the variation of the green colors are the confidence intervals. The general trend of the power allocation for actual usage follows fixed by the design and stationary allocation for CIR scenarios. The results for the other satellites are the same but shifted in time based on their initial longitude. Each of the black vertical lines separates a 6-hour orbit.

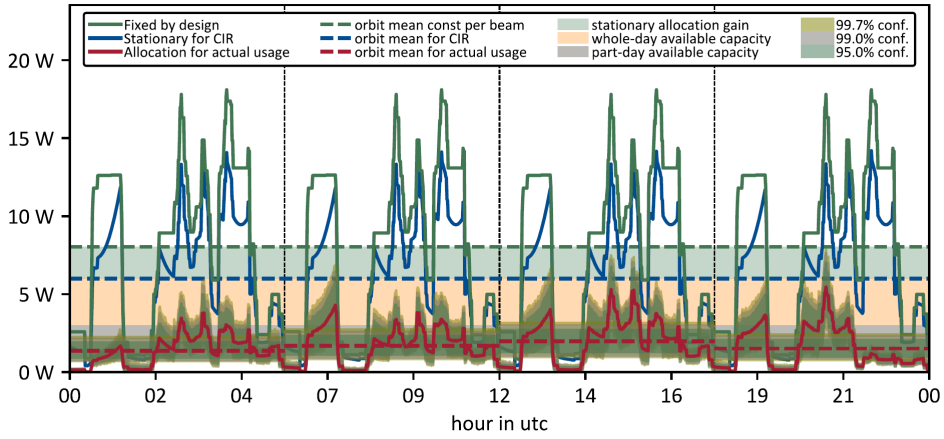

Figure 7: Power allocation for the three different scenarios. 1-minute granularity over 24 hours separated into four 6-hour orbits. The plot shows the power allocation for satellite #3.

We compute the available whole-day and part-day capacity based on a per orbit basis (to account for the smoothening effect of on-board batteries), and compute the whole-day capacity based on the mean orbital distribution for which the actual usage is the highest. Values for satellite #3 are displayed in Table 6

		Orbit 1 $t \in [0h, 6h]$		Orbit 2 $t \in [6h, 12h]$		Orbit 3 $t \in [12h, 18h]$		Orbit 4 $t \in [18h, 24h]$	
The three scenarios	**Fixed by design (capacity maximum) [W]**	8.06		8.06		8.06		8.06	
	Stationary for CIR [W]	6.01	75%	6.01	75%	6.01	75%	6.01	75%
	Dynamic for actual usage [W]	2.10	26%	2.55	32%	2.97	37%	2.28	28%
Available capacity	**Stationary allocation lift from stationary for CIR to fixed by design [ΔW]**	2.05	25%	2.05	25%	2.05	25%	2.05	25%
	Whole-day available capacity [ΔW]	3.02	38%	3.02	38%	3.02	38%	3.02	38%
	Part-day available capacity [ΔW]	0.87	11%	0.42	5%	0.00	0%	0.69	9%

Table 6: Summary of the mean orbit power consumption of the three different power allocation scenarios and delta capacities. The percentages are relative to the fixed by design allocation.

The resulting whole-day available capacity is 3.02 W at any orbit (or 38% of the maximum capacity). This number is close to the 36% obtained from the toy problem in [46]. During the orbits for which the traffic is lower, part-time available capacity becomes available. The maximum is during the first orbit with 0.87 W or 11%. Compared to the toy problem, the part-time available capacity is lower. The reason for that is that the constellation in this article is MEO instead of geostationary, and we compute the part-time capacity on a per orbit basis. Both factors have a smoothening effect on power consumption.

In sum, the results indicate that stationary allocation for CIR reduces the capacity usage by 25% (100% - 75%) compared to the fixed by design allocation, while dynamic allocation for actual traffic reduces it by 63% (100% - 37%).

6. Results

In this section we discuss the results of the simulations to translate the 38% available whole-day capacity, and up to 11% part-day capacity, computed in the previous section, into additional revenues. The section is divided into

the two main approaches presented in Section 3: *Selling more capacity through existing SLAs* and *Selling additional products*. Each subsection first discusses the results and then draws specific conclusions and limitations for each approach. Finally, a more in-depth comparison between the two analyses is provided at the end of the section.

6.1. Selling more capacity through existing SLAs

6.1.1. Discussion of the results

Figure 8 shows the monthly revenues (normalized to the baseline), the average price (normalized to the average 2019 price), and the total contracted CIR capacity. For each of these metrics, we compare the two algorithms with the baseline across the three elasticity cases. The percentage above the red bar is the relative improvement of the gradient approach compared to the heuristic.

For the unit elastic case, the monthly revenues do not change between the different algorithms. That is ordinary and validation of the simulation, since we expect that behavior from a unit elastic demand. Compared to the baseline, the gradient optimization does not make changes to the pricing. At the same time, the equal decreasing heuristic decreases average prices across all customers by 32%, resulting in 35% more sold capacity.

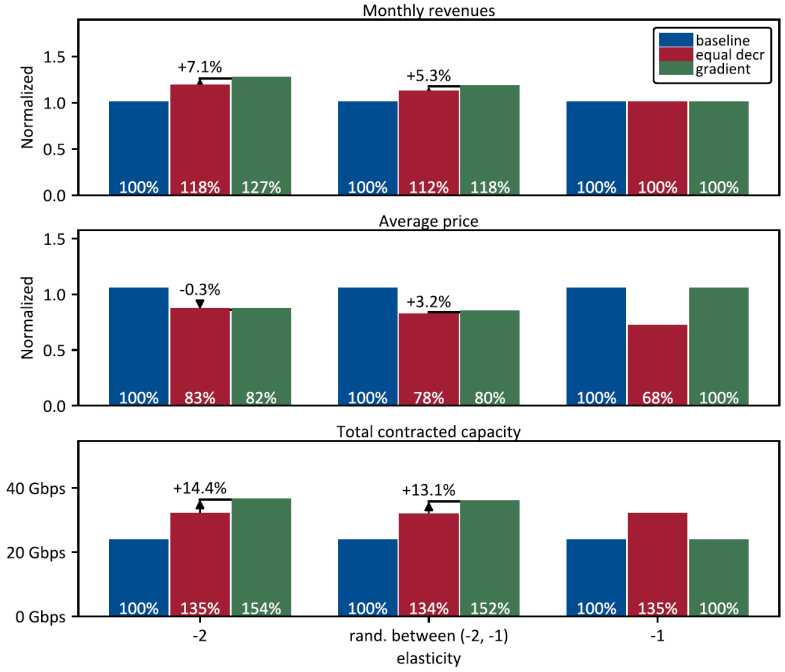

Figure 8: Monthly revenues, average price, total contracted CIR capacity for the three elasticity cases compared between the baseline and the two pricing algorithms.

6.1.2. Conclusions and limitations

With the current approach, we studied how the operator can sell the additional capacity to the same customer through the same SLA. We compared two algorithms, a heuristic, and a gradient optimization across three elasticity cases: elastic (-2), elastic range (from -2 to -1), and inelastic. From the simulation, we conclude the following main points:

- The 38% available capacity is translatable into 0-27% additional revenues depending on the elasticity. More elastic demand results in higher improvements.
- The gradient optimization outperforms the heuristic approach by 5-7%.

Besides the assumptions outlined at the beginning of Section 6, the analysis revealed some further limitations of the simulation:

- The gradient optimization is only stable for elastic demand; for inelastic demand, the optimum is to sell the smallest increment of capacity. The log-linear price elasticity function that we use has the property that the elasticity is constant at all points on the curve. Other fitting functions with varying elasticity as a function of the price might be used for better representations of the extremes.

- The assumption for the simulation that all power adjustments are sufficient limited the demand increase for a few users since more demand would require additional frequency spectrum. Including a re-computation of the frequency plan in the simulation could potentially improve the optimized revenues.

6.2. Selling additional products

6.2.1. Discussion of the results

We simulate each of the nine combinations from Table 1 with four different elasticity cases: elastic with -2, random between -2 and -1, unit elastic at -1, and inelastic with -0.5. The results are reported for the baseline as well as for the decreasing heuristic algorithm and gradient optimization.

Figure 9 displays the results in the same format as Figure 8 for one of the nine combinations with $k_{rev,fract} = 0.2$ and $k_{adj} = 0.5$.

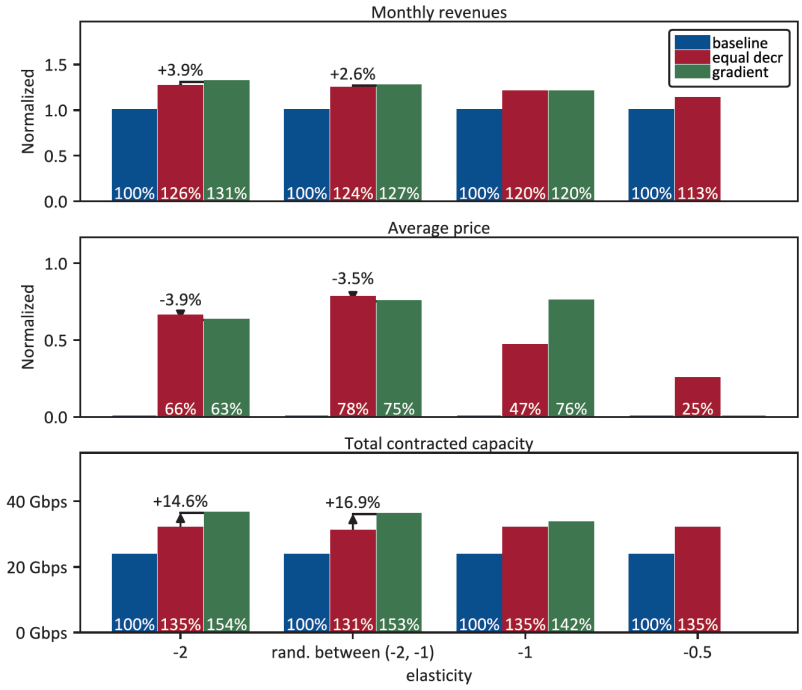

Figure 9: Results for the combination $k_{rev,fract} = 0.2$ and $k_{adj} = 0.5$.

The bar plot compares the two algorithms with the baseline, and the percentage points above the red bar point out the lift from the gradient optimization. The baseline has no average prices since the products are not part of the baseline. Therefore, we report the percentage points of the two algorithms concerning the normalized scale. Additionally, the gradient optimization does not produce consistent results for the inelastic demand for the reasons discussed in the limitations part of the previous subsection.

The revenue lift of unit elasticity case is 20% above the baseline, which validates our modeling of the parameter $k_{rev,fract}$ (0.2 for the results in the figure). For the elastic case, the revenue lift is 24-31%, while it is 13% for the inelastic -0.5 case with the heuristic. The gradient optimization achieves a 3-4% increase. Throughout the cases, the algorithms increase the total contracted capacity by 31-54%.

While Figure 9 only presents the results for one of the nine combinations, significant differences of the results from other combinations are also discussed.The case of $k_{rev,fract} = 0.1$ and $k_{adj} = 1$, with the equally increasing heuristic, yields the smallest additional revenues with 11% (just one percentage point above the baseline). The highest lift of 98% occurs for the -2 elastic case and $k_{rev,fract} = 0.5$ and $k_{adj} = 0.2$. The lift from the gradient optimization ranges between 1.2% and 8.5%.

On the one hand, for all elastic cases, smaller k_{adj} combinations increase the revenues across algorithms. Referring back to Figure 2, that behavior makes sense since more quantities can be sold at a higher price. On the other hand, larger values for k_{adj} result in more additional revenues in the inelastic case.

6.2.2. Conclusions and limitations

With this second approach, we analyzed what revenues gains are achievable by selling the available capacity through new products, such as a spot instances add-on. Given the considerable uncertainty about the pricing of these options, we explored a wide range parameterized by $k_{rev,fract}$ and k_{adj}. Based on the discussion of the results, we conclude the following points:

- The 38% available capacity is translatable into 2-102% additional revenues, which is more than the range of 0-27% we computed for the selling of the capacity through the existing Classical SLAs from the previous Section.
- Products with a higher initial price for a smaller value of Mbps perform greater if the demand is elastic. The vice versa is true for products with lower reference prices for a larger value of Mbps and inelastic demand.

We identify these limitations for the analysis presented in this Section:

- The algorithms vary prices without considering the impact on internal cannibalization. It is not an issue to include internal cannibalization in the optimization formulation. However, the critical point is to have data available on the amount of internal cannibalization as a function of price points.
- The up to 11% available part-day capacity remains unallocated for after the optimization, leaving room for additional revenues. The time-of-day pricing and Two Classes of Service are novel SLAs that can smoothen out the diurnal variations and therefore reduce part-day capacity.

6.3. Summary of the two analyses

The previous two subsections, 6.1 and 6.2 , presented the results and first conclusions of the two different approaches for monetization of the available capacity. We compared two algorithms: equally decreasing heuristic and gradient optimization to mimic different sophistication of the RM system. The reported results contain two central metrics: first, the additional revenue improvement extracted from the 38% available capacity (made available by dynamic resource allocation). Second, the lift of the more sophisticated gradient optimization RM compared to the second-best heuristic. Table 7 and Figure 10 summarize the data.

	Selling more to existing customers	
	Selling more capacity through existing SLAs	**Selling additional products**
38% available capacity is translatable into how much additional revenues	0-27%	[2-102%] #
Critical assumption		No internal cannibalization
Range depends on	• Demand for use elasticity • Sophistication of pricing algorithm	• Demand elasticity • price point for add. product • Sophistication of pricing algorithm
Separate price for	Customer	Customer
Price affects	• Demand for use	• Demand for add. product
Price does not affect	• Access demand • Market share	
Lift in revenues through more sophisticated RM over heuristic pricing (gradient optimization over second best heuristic)	5-7%*	[1-9%] #*

ranges in brackets should be taken with care due to the listed critical assumption
* ranges exclude the unit elastic case, for which the algorithm cannot influence the revenues

Table 7: Comparison of the additional revenues potential and the revenue lift from sophisticated RM across the four analyses.

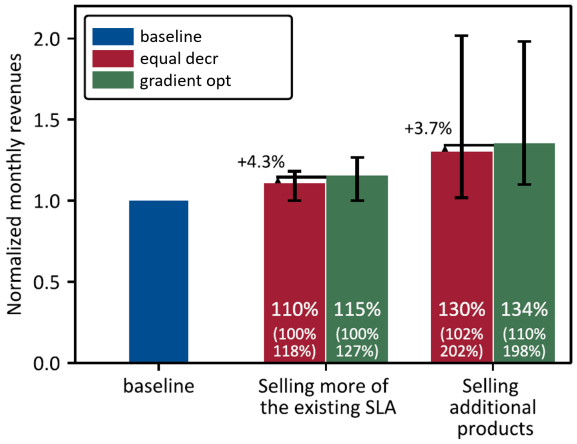

Figure 10: Comparison of the results across the three analyses. Besides the (heuristic) gradient approach, the second-best heuristic is plotted. The percentage numbers in parentheses indicate the error bar numbers, which are the minimum and maximum numbers. The bars show the means.

The wide ranges in the numbers are due to considerable variation in our assumptions about elasticities. The boundary case is a unit elastic demand, in which the additional revenues are 0% and 2% for selling more capacity and selling additional products respectively.

The selling of additional products has the most extensive upside range. However, this number should be taken with care for two reasons: first, our assumptions about the uncertainty vary the most, and hence we expect a broad range of possible revenues. Second, we did not account for the effect of internal cannibalization, which is likely to affect the revenue generated by the base Classical SLAs. If there is a considerable internal cannibalization effect, the value of additional products is smaller. We expect that cannibalization varies between products and segments and might be best estimated through sales expert knowledge. A realistic assessment of internal cannibalization would allow

for tightening the range of possible revenue gains and allowing for more nuanced conclusions regarding additional products. We believe it is a valuable future research direction, as new products could provide an effective way to prevent (or at least slow down) a potential race-to-the-bottom. It allows operators to differentiate themselves and make service offerings stickier to customers, giving the operator a first-mover advantage with novel SLAs.

To conclude this section, we switch our focus to the discussion of the second metric: the lift through more sophisticated RM techniques. We observe a lift between 1% and 9%. Note that these ranges do not include the unit elastic case since the algorithms do not influence the revenues. However, the average lift numbers shown in Figure 10 include this case. The interpretation of these numbers should be how efficient the pricing principles are with the available capacity. Our intention with the equally decreasing heuristic is to approximate the current pricing policies operators might follow. Comparing to this baseline, a more sophisticated RM system can then use the capacity more efficiently, and hence, increase the revenues by 1-9%.

7. Conclusions

To summarize, in this paper we brought the pieces of the satcom RM framework together and analyzed its value under different assumptions. We outlined a pricing optimization approach that works with a computationally expensive evaluation function, such as our resource management part of the framework. We set up a simulation that contains actual data and used this as a baseline throughout the paper. From there, we first quantified the benefit of dynamic power allocation to be 38% whole-day available capacity. Second, we looked at two different scenarios of monetizing this 38%: selling more of existing SLAs and selling additional products.

The results let us draw the following conlcusions:

- Stationary CIR allocation reduces capacity usage by 25% compared to fixed by design allocation. Dynamic allocation for actual traffic reduces it by 63% and frees up 38% whole-day available capacity compared to the stationary CIR allocation.
- The 38% available capacity is translatable into 0-27% of additional revenues, mainly depending on customers' price elasticity.
- Monetizing the available capacity through additional products has the highest potential but also the most significant uncertainties. It comes with the risk of internal cannibalization.
- More sophisticated pricing algorithm within the RM framework lift revenues between 1-9% over the best heuristic pricing approach. This benefit is sustainable if the market has limited price transparency or competitors' satellite systems are different.
- Satcom RM is not a zero-sum game since it improves revenues by using the capacity more efficiently.

Appendix A. Binary search algorithm to find the element where available capacity is zero

Formally, we define the ordered list L with the length N_l where p_j is the price for user j.

$$L = \{l_i = (p_j \; \forall j \in 1...N_u)\} \; \forall i \in \{1...N_l\} \tag{A.1}$$

With that definition, we outline the pseudo-code for a recursive binary search algorithm [47] in Algorithm 1. Line 4 is computationally the most expensive part of the algorithm since it involves a rerun of the resource allocation process.

Algorithm 1 Pseudo code for the recursive binary search algorithm for pricing optimization

Input: $\mathcal{L}$ with declining prices, i.e., increasing capacity
Output: l

def $binarySearch(\mathcal{L}, lb, ub)$: // function definition with lower and upper bound (lb, ub)
  $mid = 1 + (ub - lb) // 2$ // floor division to obtain the middle index
  if $ub \geq lb$ **do**
    $C_{avail} = availableCapacity(\mathcal{L}[mid])$ // compute available capacity for set of prices
    if $C_{avail} == 0$ **do** // if available capacity zero, return current prices
      return $\mathcal{L}[mid]$
    else if $C_{avail} > 0$ **do** // if available capacity larger zero, search to the right
      return $binarySearch(\mathcal{L}, mid, ub)$
    else do // if available capacity smaller zero, search to the left
      return $binarySearch(\mathcal{L}, lb, mid)$
    end if
  else
    return $\mathcal{L}[mid]$
  end if

References

[1] Gunter, ipstar 1 (thaicom 4, measat 5, synertone 1), accessed: March 2019.
URL *https : //space.skyrocket.de/doc_sdat/ipstar − 1.htm*

[2] SES, O3b mpower, accessed: March 2019.
URL *https : //www.ses.com/networks/networks − and − platforms/o3b − mpower*

[3] Gunter, O3b mpower 1, ..., 11 (o3b 21, ..., 31), accessed: March 2019.
URL *https : //space.skyrocket.de/doc_sdat/o3b − 21.htm*

[4] G. Buttazzoni, M. Comisso, A. Cuttin, M. Fragiacomo, R. Vescovo, R. Vincenti Gatti, Reconfigurable phased antenna array for extending cubesat operations to ka-band: Design and feasibility, Acta Astronautica 137 (2017) 114 – 121. doi:https://doi.org/10.1016/j.actaastro.2017.04.012.
URL `http://www.sciencedirect.com/science/article/pii/S0094576517302461`

[5] A. Canabal, R. Jedlicka, A. Pino, Multifunctional phased array antenna design for satellite tracking, Acta Astronautica 57 (12) (2005) 887 – 900. doi:https://doi.org/10.1016/j.actaastro.2005.03.072.
URL `http://www.sciencedirect.com/science/article/pii/S0094576505001955`

[6] A. I. Aravanis, B. S. MR, P.-D. Arapoglou, G. Danoy, P. G. Cottis, B. Ottersten, Power allocation in multibeam satellite systems: A two-stage multi-objective optimization, IEEE Transactions on Wireless Communications 14 (6) (2015) 3171–3182.

[7] A. Paris, I. del Portillo, B. G. Cameron, E. F. Crawley, A genetic algorithm for joint power and bandwidth allocation in multibeam satellite systems, in: 2019 IEEE Aerospace Conference, IEEE, 2019, pp. 1–15.

[8] Y. He, Y. Jia, 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 13 (8) (2017).

[9] N. Pachler, J. J. G. Luis, M. Guerster, E. F. Crawley, B. G. Cameron, Allocating power and bandwidth in multibeam satellite systems using particle swarm optimization, in: 2020 IEEE Aerospace Conference, 2020.

[10] F. R. Durand, T. Abrão, Power allocation in multibeam satellites based on particle swarm optimization, AEU - International Journal of Electronics and Communications 78 (2017) 124–133. doi:https://doi.org/10.1016/j.aeue.2017.05.012.
URL `https://www.sciencedirect.com/science/article/pii/S1434841116303739`

[11] H. Wang, A. Liu, X. Pan, J. Yang, Optimization of power allocation for multiusers in multi-spot-beam satellite communication systems, Mathematical Problems in Engineering 2014 (2014).

[12] J. J. G. Luis, M. Guerster, I. del Portillo, E. F. Crawley, B. G. Cameron, Deep reinforcement learning architecture for continuous power allocation in high throughput satellites, in: Reinforcement Learning for Real Life Workshop at 2019 International Conference on Machine Learning, 2019.

[13] G. Cocco, T. de Cola, M. Angelone, Z. Katona, S. Erl, Radio resource management optimization of flexible satellite payloads for dvb-s2 systems, IEEE Transactions on Broadcasting 64 (2) (2018) 266–280. doi:10.1109/TBC.2017.2755263.

[14] X. Hu, S. Liu, R. Chen, W. Wang, C. Wang, A deep reinforcement learning-based framework for dynamic resource allocation in multibeam satellite systems, IEEE Communications Letters 22 (8) (2018) 1612–1615. doi:10.1109/LCOMM.2018.2844243.

[15] S. Liu, X. Hu, W. Wang, Deep reinforcement learning based dynamic channel allocation algorithm in multibeam satellite systems, IEEE Access 6 (2018) 15733–15742. doi:10.1109/ACCESS.2018.2809581.

[16] J. J. G. Luis, E. F. Crawley, B. G. Cameron, Applicability and challenges of deep reinforcement learning for satellite frequency plan design, in: 2021 IEEE Aerospace Conference, 2021.

[17] U. Park, H. W. Kim, D. S. Oh, B.-J. Ku, A dynamic bandwidth allocation scheme for a multi-spot-beam satellite system, ETRI Journal 34 (4)

(2012) 613–616. doi:https://doi.org/10.4218/etrij.12.0211.0437.
URL https://onlinelibrary.wiley.com/doi/abs/10.4218/etrij.12.0211.0437
[18] J. Camino, C. Artigues, L. Houssin, S. Mourgues, Mixed-integer linear programming for multibeam satellite systems design: Application to the beam layout optimization, in: 2016 Annual IEEE Systems Conference (SysCon), 2016, pp. 1–6. doi:10.1109/SYSCON.2016.7490613.
[19] A. Kyrgiazos, Irregular beam sizes and non-uniform bandwidth allocation in HTS, 2013. doi:10.2514/6.2013-5617.
[20] N. Pachler, M. Guerster, I. del Portillo, E. F. Crawley, B. G. Cameron, Static beam placement and frequency plan algorithms for leo constellations, International Journal of Satellite Communications and Networking (2020).
[21] J. Liang, Y. he Zhu, Y. zhong Luo, J. cheng Zhang, H. Zhu, A precedence-rule-based heuristic for satellite onboard activity planning, Acta Astronautica 178 (2021) 757 – 772. doi:https://doi.org/10.1016/j.actaastro.2020.10.020.
URL http://www.sciencedirect.com/science/article/pii/S0094576520306159
[22] B. Gaudet, R. Linares, R. Furfaro, Terminal adaptive guidance via reinforcement meta-learning: Applications to autonomous asteroid close-proximity operations, Acta Astronautica 171 (2020) 1 – 13. doi:https://doi.org/10.1016/j.actaastro.2020.02.036.
URL http://www.sciencedirect.com/science/article/pii/S0094576520301090
[23] J. H. Park, H. Jung, C. H. Lim, T. Chang, The economic impact analysis of satellite development and its application in korea, Acta Astronautica 177 (2020) 9 – 14. doi:https://doi.org/10.1016/j.actaastro.2020.06.031.
URL http://www.sciencedirect.com/science/article/pii/S0094576520303982
[24] L. del Monte, L. Scatteia, A socio-economic impact assessment of the european launcher sector, Acta Astronautica 137 (2017) 482 – 489. doi:https://doi.org/10.1016/j.actaastro.2017.01.005.
URL http://www.sciencedirect.com/science/article/pii/S0094576515302599
[25] S. Yi, H.-J. Jang, H. S. Lee, J.-P. Yu, S. Kim, J. Lee, H.-Y. Hur, Economic value analysis of the return from the korean astronaut program and the science culture diffusion activity in korea, Acta Astronautica 87 (2013) 1 – 7. doi:https://doi.org/10.1016/j.actaastro.2012.12.010.
URL http://www.sciencedirect.com/science/article/pii/S0094576512004900
[26] NSR, Vsat and broadband satellite markets, 17th edition, accessed: November 2019 (2019).
URL https://www.nsr.com/research/vsat-broadband-satellite-markets-19th-edition/
[27] M. Stanley, Space: Investing in the final frontier (2017).
URL http://www.fullertreacymoney.com/system/data/files/PDFs/2017/October/20th/msspace.pdf
[28] A. Szalay, J. Gray, 2020 computing: Science in an exponential world, Nature 440 (2006) 413–4. doi:10.1038/440413a.
[29] K. Kambatla, G. Kollias, V. Kumar, A. Grama, Trends in big data analytics, Journal of Parallel and Distributed Computing 74 (7) (2014) 2561–2573, special Issue on Perspectives on Parallel and Distributed Processing. doi:https://doi.org/10.1016/j.jpdc.2014.01.003.
URL https://www.sciencedirect.com/science/article/pii/S0743731514000057
[30] M. Chen, S. Mao, Y. Liu, Big data: A survey, Mobile Networks and Applications 19 (2) (2014) 171–209. doi:10.1007/s11036-013-0489-0.
URL https://doi.org/10.1007/s11036-013-0489-0
[31] J. Gantz, D. Reinsel, Extracting value from chaos, in: IDC iview, Vol. 1142, 2011.
[32] E. Team, The exponential growth of data (2017).
URL *https* : //*insidebigdata.com*/2017/02/16/*the* − *exponential* − *growth* − *of* − *data*/
[33] SES, Unleashing the potential of an empowered - world with the launch of o3b mpower, (2017).
URL *https* : //*www.ses.com*/*sites*/*default*/*files*/2017 − 09/170908$_S$*ES*%20*Launches*%20*O*3*b*%20*mPOWER*$_F$*INAL.pdf*
[34] S. L. Kota, Broadband satellite networks: trends and challenges, in: IEEE Wireless Communications and Networking Conference, 2005, Vol. 3, 2005, pp. 1472–1478 Vol. 3. doi:10.1109/WCNC.2005.1424732.
[35] J. Farserotu, R. Prasad, A survey of future broadband multimedia satellite systems, issues and trends, IEEE Communications Magazine 38 (6) (2000) 128–133. doi:10.1109/35.846084.
[36] N. Hosseini, H. Jamal, J. Haque, T. Magesacher, D. W. Matolak, Uav command and control, navigation and surveillance: A review of potential 5g and satellite systems, in: 2019 IEEE Aerospace Conference, 2019, pp. 1–10. doi:10.1109/AERO.2019.8741719.
[37] K. Talluri, G. van Ryzin, The Theory and Practice of Revenue Management, Kluwer Academic Publishers, Norwell, MA, USA, 2004.
[38] M. Aldebert, M. Ivaldi, C. Roucolle, Telecommunications demand and pricing structure: An econometric analysis, Telecommunication Systems 25 (1) (2004) 89–115. doi:10.1023/B:TELS.0000011198.50511.a4.
URL https://doi.org/10.1023/B:TELS.0000011198.50511.a4
[39] M. Aldebert, M. Ivaldi, C. Roucolle, Telecommunications demand and pricing structure: An econometric analysis, in: Telecommunication Systems, 2004, pp. 89–115.
[40] R. E. Park, B. M. Wetzel, B. M. Mitchell, Price elasticities for local telephone calls, Econometrica 51 (6) (1983) 1699–1730.
URL http://www.jstor.org/stable/1912113
[41] F. A. Wolak, Can universal service survive in a competitive telecommunications environment? evidence from the united states consumer expenditure survey, Information Economics and Policy 8 (3) (1996) 163–203. doi:https://doi.org/10.1016/0167-6245(96)00009-1.
URL https://www.sciencedirect.com/science/article/pii/0167624596000091
[42] J. P. Gatto, H. H. Kelejian, S. W. Stephan, Stochastic generalizations of demand systems with an application to telecommunications, Information Economics and Policy 3 (4) (1988) 283–309.
[43] T. Garín-Muñoz, T. Perez-Amaral, Econometric modelling of spanish very long distance international calling, Information Economics and Policy 10 (2) (1998) 237–252.
[44] H. Ouwersloot, P. Rietveld, On the distance dependence of the price elasticity of telecommunications demand; review, analysis, and alternative theoretical backgrounds, The annals of regional science 35 (4) (2001) 577–594.
[45] P. Hackl, A. H. Westlund, Demand for international telecommunication time-varying price elasticity, Journal of Econometrics 70 (1) (1996) 243–260.
[46] M. Guerster, J. Grotz, P. Belobaba, E. F. Crawley, B. G. Cameron, Revenue management for communication satellite operators – opportunities and challenges, In 2020 IEEE Aerospace Conference (2020).
[47] T. H. Cormen, C. E. Leiserson, R. L. Rivest, C. Stein, Introduction to algorithms, MIT press, 2009.