Problem representation of dynamic resource allocation for flexible high throughput satellites

Markus Guerster, Juan Jose Garau Luis, Edward Crawley, Bruce Cameron
Massachusetts Institute of Technology
77 Massachusetts Ave 33-409
Cambridge, MA 02139
857-999-6103
{guerster, garau, crawley, bcameron}@mit.edu

Abstract— **Within the next years, flexible high-throughput (HT) satellite with 100s-1000s of beams will be launched to provide broadband connectivity to a variety of customers on Earth. The user demand, especially in the mobility sector, is expected to have large diurnal variations. To follow this dynamic demand behavior, many HT satellites will be equipped with flexible power and bandwidth capabilities. This flexibility comes with a large number of adjustable parameters. Optimization of these parameters ensures that the demand can be met with the minimum required resources, resulting in an efficient utilization of assets on orbit. The challenge lies in the high dimensionality of the problem, where manual resource allocation will quickly become impractical. This paper develops a representation of the dynamic resource allocation problem and outlines an approach to solve the problem including different sets of algorithms such as deterministic solvers and heuristic approaches from artificial intelligence.**

TABLE OF CONTENTS

1. INTRODUCTION

The space industry is increasingly interested in providing broadband internet access through satellites. The Federal Communications Comission (FCC) received 11 applications from commerical companies for new high-throughput (HT) non-geostationary satellites [1], especially large LEO constellations, such as from Telesat [2], OneWeb [3] and SpaceX [4]. In addition to those newly proposed mega-constellations, established players in the communication market expect to launch MEO and GEO satellites to specifically provide broadband connectivity. Viasat is planning to launch Viasat-3 to provide 100+ Mbps broadband access by 2020 [5]. SES is launching a fleet of mPower spacecrafts to supplement their O3b MEO satellites within the next years [6], as well as SES-17 – a HT GEO satellite that will provide connectivity to America and the Atlantic Ocean [7].

The idea of providing broadband internet access through space is not new. In the 90s, companies such as Iridium, Globalstar, and Orbcomm had similar ideas, but market uptake was poor [8-10]. However, over the past two decades, technological developments such as digital communication payloads, advanced modulation, multi-beam antennas, and advanced manufacturing are used by the new generation of communication satellites. These advances lead to performance increases (in terms of data throughput) while at the same time reducing capital and operational costs. This cost reductions could increase the competitiveness of satellite broadband access relative to terrestrial broadband solutions. With the expected growth of demand, satellite broadband internet is expected to grow as well and might capture a larger market share in the coming years [11]. This is especially true for the mobility sector (aircraft and ships), where terrestrial alternatives perform more poorly [12].

For current analog bent-pipe satellites, the spatial distribution of capacity is fixed early in the design process based on demand predictions. If during the lifetime of the satellite, the demand behaves differently than predicted, the configuration of the satellite cannot be changed to reflect the shift [13]. To avoid this lock-in, more new satellites use flexible communication payloads in the form of digital processors, multi-beam antennas, and flexible amplifiers capable of adapting to changing demand [14]. While these technologies are expected to provide great benefit, the operation of these satellites needs to be rethought. The degrees of freedom are orders of magnitude above those of current classical analog bent-pipe satellites. Current resource allocation of power and bandwidth is mainly done manually. While this is still a workable approach for satellites with limited degrees of freedom, a new autonomous dynamic resource management (DRM) needs to be developed to manage future satellites with significantly higher degrees of freedom.

Our objective with this paper is to present our first steps in developing a DRM for the future generation of highly flexible HT satellites.

The need for dynamic resource management is not new. Several authors addressed this issue by developing optimization approaches [15-19]. Specifically, Sharma [20] and Lagunas [21] developed an approach to monitor the spectrum and allocate bandwidth and exploit beamforming capabilities.

978-1-5386-6854-2/19/$31.00

However, little work has been done to develop a descriptive representation of the requirements for such a dynamic resource management (DRM) tool. The cited literature addresses technical sub-problems of the DRM. With this paper, our specific objective is to provide a more holistic system view. We will decompose the DRM into its main functionalities (Section 2), describe the behavior for perfect and imperfect information (Section 3), followed by a discussion of the different states and configuration of the DRM (Section 4). We discuss the limitations of the chosen representation in Section 5 and conclude the paper in Section 6 with an outlook on algorithms that are worth exploring for the functional blocks within the DRM.

2. PROBLEM DECOMPOSITION

The main task of the DRM tool is to control satellites in an optimal way, so resources are used efficiently, and the data rates provided to users can be maximized. To do so, the current operational infrastructure needs to be supplemented by additional functional blocks.

Loosely speaking, the current infrastructure consists mainly of the satellite *physical network,* which provides telemetry to a human operator and is controlled by telecommands. On a high-level, we decompose the *physical network system* into 1…N satellites with gateways and user terminals for a variety of use cases. Specifically, we assume that future satellites have built-in flexibility in one or more of the following areas that make traditional manual operation impractical:

- RF power-level adjustment per beam
- Bandwidth adjustment per beam
- Beamforming and pointing (not considered in this problem description)

To operate satellites with those high degrees of freedom, additional functional blocks are required, and the combination of these blocks results in the *DRM architecture*. Each main functionality is depicted in its own block in Figure 1 and described below.

Virtual models and constraints (VMC) to simulate the physical network

The virtual models and constraints (VMC) are a virtual copy of the physical network which simulates the telemetry and telecommands data streams. Both, the VMC and physical network connect to either an offline simulation or real-time operation instance of the Real Time Engine (RTE), respectively.

Demand estimator to simulate the real demand

The demand estimator is continuously trained with real or simulated data sets to predict a pool of potential users throughout time that is a simulation of the real demand the operational system might see.

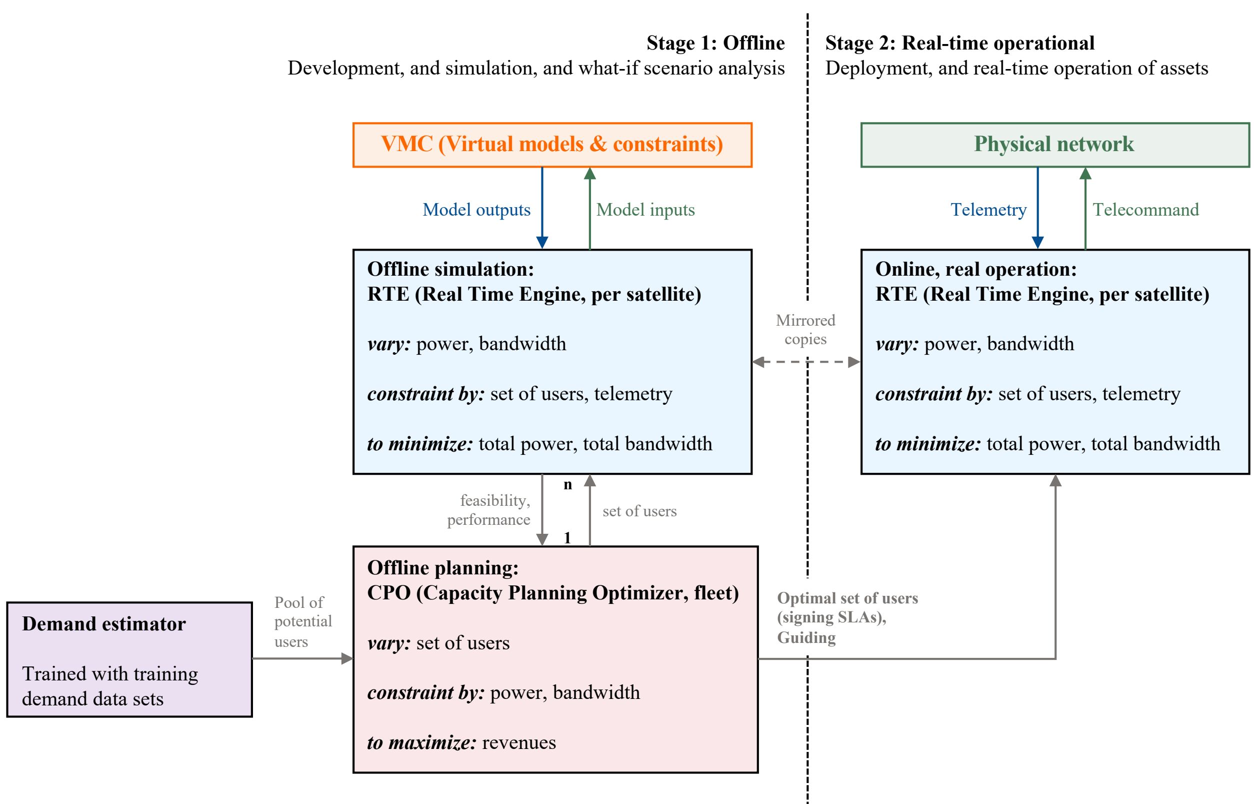

Figure 1: Architecture of the dynamic resource allocation (DRM). Decomposing the optimization functionality into capacity planning optimizer (CPO) and an offline and real-time instance of the real time engine (RTE)

Real Time Engine (RTE) to control the physical network and VMC

The RTE is responsible for assigning optimal power and bandwidth levels to satisfy all demand using the minimum amount of resources possible and considering all the constraints introduced by the physical network or VMC, respectively. The RTE is continuously working in cycles (real-time operating system). It determines which resources need to be updated at every iteration. The RTE is governed by the RTE algorithm.

Capacity planning optimizer (CPO) to guide the RTE

The CPO suggests for which set of potential customers a new service level agreement (SLA) should be signed. When this optimal set of new customers is incorporated into the existing pool of customers, monetary measures are increased without reaching saturation of the satellite. The CPO is an offline system and compares different sets of potential users (with different estimated demands). It runs an offline instance of the RTE algorithm. This way, it guarantees the real-time RTE will be able to handle the different demand scenarios involved in a hypothetical new complete set of users (current + subset of the potential). The CPO is governed by the CPO algorithm.

To summarize, the CPO guidance and the RTE has control functionality with their specific algorithms. Due to the high dimensionality and hardness of the problems, we do not expect that they necessarily find the optimal solution. Therefore, we can say that the CPO and RTE determine the *optimality* of the DRM.

3. PERFECT VERSUS IMPERFECT INFORMATION

The VMC aims to model reality and the demand estimator predicts demand behavior. No model has perfect information about reality and therefore the quality of the VMC and demand estimator significantly determines the *accuracy* of the DRM and the *conservatism* the CPO must have. Since the level of information is a significant driver, we explain the behavior of the DRM for two scenarios: perfect and imperfect information.

Perfect information

Assumptions:

- The DRM has perfect information about the variation of the demand, and there is no uncertainty in these predictions.
- The VMC are a perfect representation of the physical network system. All contingencies scenarios can be captured and predicted perfectly by the model.

Process:

- The CPO finds the optimal set of users by executing an offline instance of the RTE which runs the VMC to simulate the real system behavior.
- Once the optimal set of users is found, the CPO passes on the corresponding optimal plan (guidance) to the online RTE instance.
- The online RTE operates the system based on this optimal plan. The online RTE does not need to deviate from the optimal plan since the demand predictions are perfect, the VMC is a perfect copy of the physical network, and all contingency scenarios are captured perfectly by the model.

Imperfect information

Assumptions:

- The DRM has imperfect information about the variation of the demand, and there is considerable uncertainty in these predictions.
- The virtual models & constraints (VMC) is not a representation of the physical network. Not all contingencies scenarios can be captured and predicted perfectly by the model.

Process:

- The CPO finds the optimal set of users by executing an offline instance of the RTE which runs the VMC to simulate the real system behavior. Due to an imperfect model, the CPO needs to have a certain level of conservatism (= margin).
- Once the optimal set of users is found, the CPO passes on the corresponding optimal plan (guidance) to the online RTE (even though this plan is optimal and perfect for the virtual world, it is imperfect for the real physical system due to the imperfect information).
- The online RTE operates the system based on this optimal plan. The online RTE needs to deviate from the optimal plan and needs to make real-time decisions since the demand predictions are imperfect, the VMC is an imperfect copy of the physical system, and not all contingency scenarios are captured perfectly by the model. If the conservatism of the CPO was sufficient to capture the inaccuracies of the VMC and demand estimator, the RTE can control the physical network in a way that demand is met at all times.

4. OPERATIONAL STATES OF THE DRM TOOL

After we decomposed the DRM tool, in this Section we define the different operational environments of the tool.

Depending on the ratio of satisfied demand to system capacity, the DRM tool operates in different *states* (see Figure 3). We define a state as points where the tool operates similarly. If the state changes, the DRM needs to operate significantly differently. Furthermore, we define *configurations* within a state (specific allocation of the satellite resources, assuming certain power levels and bandwidth usages for each beam). Configurations describe specific operational points within a state. The general behavior does not change if the tool moves from one configuration to the next. As a third definition, we refer with *scenario* to a specific distribution of users, whose demand may vary according to daily or other periodic patterns.

In order to visualize the different states and configuration, we introduce the demand-supply diagram shown in Figure 2. The served demand (= current users, not the potential demand) is shown on the horizontal axis, whereas the vertical axis represents the supply as throughput. The diagonal black line shows the desirable operational line where all of the demand is met at all times.

The system has an *overall capacity limit*, which defines the maximum achieved data throughput and the maximum met demand (black dotted line). However, the system can only reach the capacity limit under certain "optimal" user distributions (a user distribution = scenario). This is a perfectly distributed and infinite demand. Under this scenario, the satellite has a full utilization of power and bandwidth resources on all beams.

When user distribution deviates from the "optimal" scenario – e.g. a large number of users clustered on a specific beam's footprint – the system is not limited by the overall maximum capacity but by the *maximum data throughput* the new scenario imposes. This new limit is represented with the green dotted line, whose exact position is unknown due to the high computational power required to determine it.

Since the availability of computational resources is a constraint, all algorithms under consideration are defined by their non-exhaustive solution exploration policies, i.e. algorithms that find close-to-optimal solutions in admissible time intervals. However, the use of non-exhaustive techniques materializes a new limit for the system. This new limit is the *maximum achieved performance* represented by the red dotted line in Figure 2 and accounts for the mentioned algorithm suboptimalities. If the DRM tool had an unlimited amount of computational power, the Resource Allocation algorithm could perfectly determine the optimal solution under any scenario and hence the red and green lines would always coincide.

The objective of the DRM is to improve current *State of the Art* manual allocation approaches (SoA, blue dotted line). Reflecting this goal on Figure 2, our focus is to approach the red line to the green line while distancing from the blue line. Focusing on pushing the algorithm limit up allows the system to meet greater demands with configurations lying on the operational line. When the system faces a different demand, the operational points need to be changed to a new configuration. This is represented in Figure 2, in which the operability of different configurations is shown.

Configuration 1 is well below the algorithm limit of 100% and has a mean operational point shown by the blue cross. Since the demand varies from this mean point (shown with the blue distribution), the Real-time Engine (RTE) needs to follow the demand. The demand distribution is mainly caused by diurnal, weekly or seasonal changes. Some of these contributions might be more or less predictable by the demand estimator.

Configuration 2 results if a new service level agreement (SLA) is contracted, the demand increases, and the Capacity Planning Optimizer (CPO) needs to find an optimal set of customers with a new optimal mean operational point. The variation at this point might have a different shape. It might cover a larger or smaller range, requiring from the RTE to

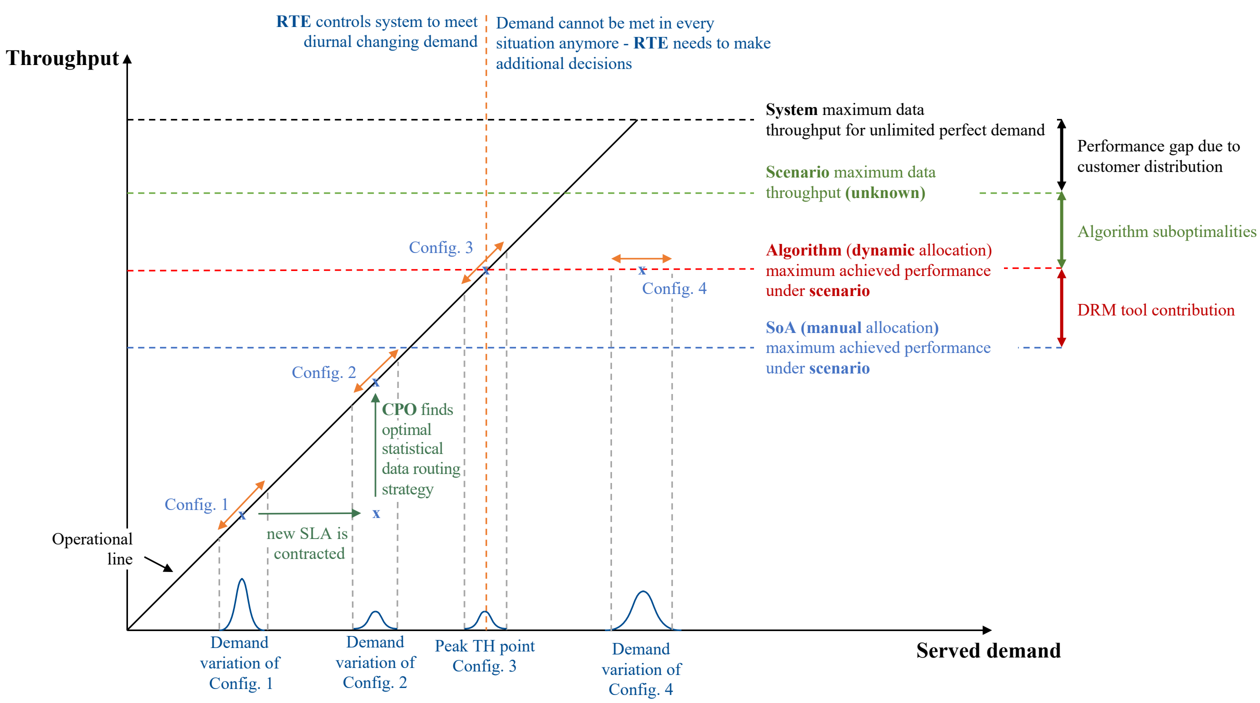

Figure 2: Demand-supply diagram to describe the different configurations

control the system over a larger or smaller operational range (compared to configuration 1).

Configuration 3 is where the demand is close to 100% of the sold capacity. In this configuration, if the demand is below 100% of the max achievable capacity, the RTE operates like in the previous configuration and controls the system to follow the demand. But if the demand exceeds the available capacity (to the right of the 100% demand vertical line), the RTE has to make additional decisions since it cannot satisfy all of the demand at all times. Possible decisions could include the reduction of data rate for lower priority users with lower contracted availability or a uniformly reduced data rate.

Configuration 4 is shown to be at the maximum achieved data rate imposed by the suboptimalities of the algorithm under the scenario, with a contracted demand largely exceeding the 100%. In this case, it is impossible to satisfy all the demand and the system is forced to either reduce data rate or drop users. In any case, this configuration involves the violation of one or more SLAs if operated for longer periods.

The discussed *conservatism* for the CPO should be chosen in a way that operation to the right of the 100% demand vertical line should be avoided.

Our experience is that diurnal variations in the demand are significant with values of 50 % not uncommon [22, 23]. Transferring this to the demand-supply diagram in Figure 2, this would mean that if more than 50 % of the available capacity is contracted, there is a probability of greater zero, that the demand will exceed the capacity at some point in time. In these cases, the RTE has to be equipped with additional decision rules to react accordingly as discussed for configuration 3. This summed variation of the complete demand can be reduced if the CPO picks a more diverse set of users.

Considering the limit imposed by the algorithm's suboptimalities and the different configurations studied, three system states are defined. A state of the system as defined can be seen as the average ability of the system to meet the demand (see Figure 3).

State 1 is where configurations 1 and 2 from Figure 2 lay. In this state, the system is always capable of meeting all the demand and its main priority is to optimize for power and bandwidth. Conservative strategies focus on always operate in this state.

State 2 is a narrow operation area in which the statistical average of the demand lays closer to the 100%. In this state, the system is constantly on the verge of the saturation point and is utilizing almost all resources possible under the assumed scenario. Such utilization does not allow for much power and bandwidth optimization. Instead, the system focuses on optimizing data rate for all users. There is a significant probability that the demand will exceed the 100% and the system will provide less data rate than the committed for some of the users.

State 3 is a non-operable state in which all of the demand is never met. Ideally, we would never want to be in that state, as even optimizing for data rate is not enough to serve all users. Being in this state implies dropping one or more users and violating their corresponding SLAs. The system might encounter itself in this state mainly due to contingency scenarios. Best practice in this state might be to optimize for the largest amount of served users possible.

To summarize, we can divide the operational environment of the system into three different states. The first state is where the demand is always below the maximum capacity, meaning the probability is zero that the demand exceeds the supply (including demand variations and uncertainty in demand

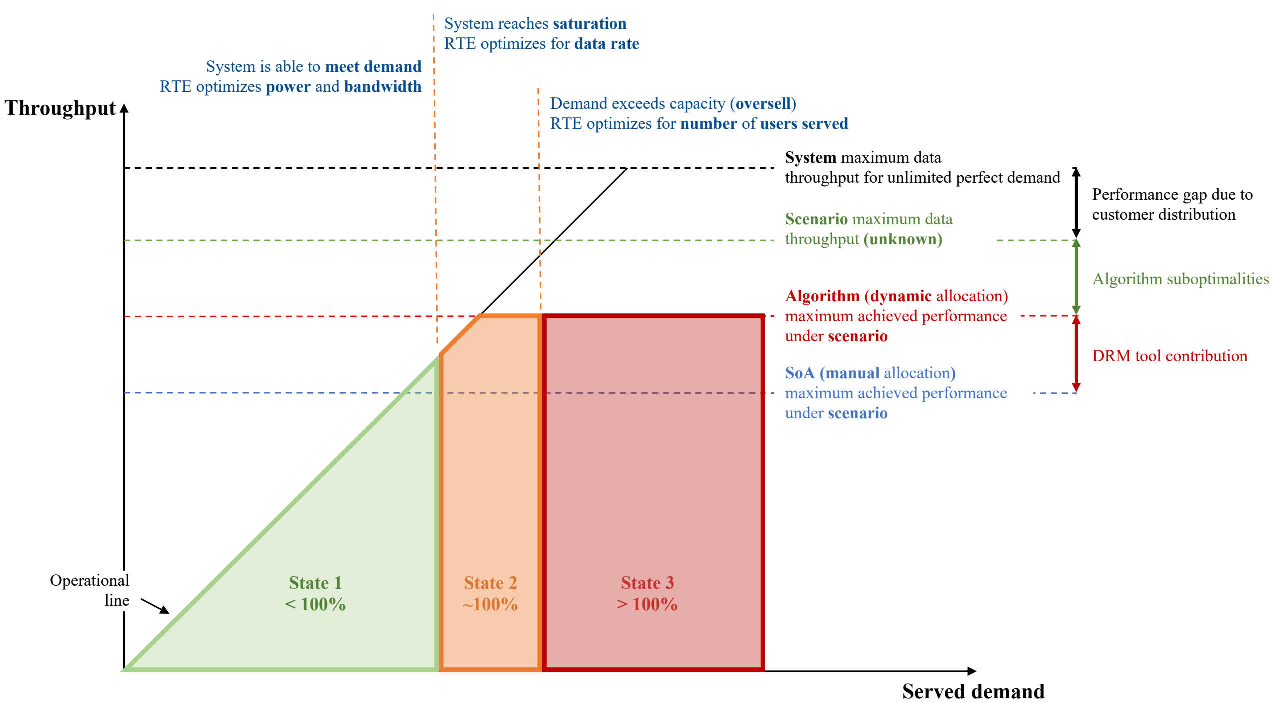

Figure 3: Demand-supply diagram to describe the different states

predictions). The second state is where the probability is greater zero that the demand will exceed the maximum capacity at certain points. The third state is where the demand always exceeds the capacity and the system is always found at the saturation point.

5. LIMITATIONS

While the presented description provides a holistic view of the DRM to capture the high-level behavior, the representation has limitations:

- The demand-supply diagram in Figure 2 and Figure 3 only represents the served users and throughput for a single system. This system can be on channel, beam, amplifier, satellite, fleet, or operator level. Additionally, the diagram does not depict information about the potential additional pool of users.

- While the states in Figure 3 are well characterized, it can be challenging in a real operational environment to clearly distinguish between states. Also, since the RTE has different objectives in each state, we have yet to determine how to manage all of the state transitions.

- To implement the DRM tool, many data cleaning steps are necessary. Additionally, the telecommands from the RTE to the physical network needs to undergo several consistency checks. These are not specifically depicted in Figure 1.

- Currently, the CPO is missing a validation step, that checks whether the combination of users in fact represents a feasible combination.

- While the DRM is supposed to operate without human interaction in the normal operational environment, for contingency cases, the human has to be included in the loop.

6. CONCLUSION AND OUTLOOK

With this paper, we presented the first steps in developing a DRM tool for future flexible high-throughput satellites. We approached the problem by first developing a holistic view of a decomposition into five functional blocks. Then, the various operational points of the DRM are defined in a supply-demand diagram, followed by a discussion of the limitations of the chosen representation.

Specifically, we found that a Real-Time Engine (RTE) needs to cycle in real-time to actively control the physical network. Depending on the demand over capacity ratio, the RTE has different objectives that potentially require different algorithm or models to operate in real-time. The Capacity Planning Optimizer (CPO) is an offline optimizer and planner that has softer time constraints and therefore can explore the solution space more exhaustively. Both, the RTE and CPO define the *optimality* of the DRM tool. For these methods to work, there are two models needed: a demand estimator and the virtual model and constraints (VMC) that represents the physical network. They both define the *accuracy* of the DRM. Both measurements become increasingly important the closer the served demand gets to the capacity limit. Approaching this high utilization of the system, we expect the DRM to have the greatest value.

Despite the limitations, we found that choosing an initial representation is a valuable step in the framing of the problem. Especially, the decomposition of the DRM into the five functional blocks and specific description of these blocks fundamentally helped us to understand the complexity of such a tool.

There are two areas of future work:

RTE and CPO

Exploring the suitability of various sets of algorithms for the RTE and CPO. This includes Genetic Algorithm, Simulated Annealing, Greedy Search, (Deep) Reinforcement Learning, Neutral Networks, and Hybrid algorithms.

Demand estimator and VMC

Developing a demand estimator that learns from real data and predicts future demand, as well as a fast and accurate VMC. Quantifying the accuracy of the demand estimator and VMC as a key to estimate the necessary conservatism of the CPO.

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] I. d. Portillo, B. G. Cameron, and E. F. Crawley, "A Technical Comparison of Three Low Earth Orbit Satellite Constellation Systems to Provide Global Broadband," presented at the 69th International Astronautical Congress (IAC), Bremen, Germany, 1-5 October, 2018, 2018.

[2] T. Canada, "Telesat Ka-band NGSO constellation FCC filing SAT-PDR-20161115-00108," 2018, Available: http://licensing.fcc.gov/myibfs/forwardtopublictabaction.do?file_number=SATPDR2016111500108.

[3] W. S. Limited, "OneWeb Ka-band NGSO constellation FCC filing SAT-LOI-20160428-00041," 2018, Available: http://licensing.fcc.gov/myibfs/forwardtopublictabaction.do?file_number=SATLOI2016042800041.

[4] S. E. Holdings, "LLC, SpaceX Ka-band NGSO constellation FCC filing SAT-LOA-20161115-00118," 2018, Available: http://licensing.fcc.gov/myibfs/forwardtopublictabaction.do?file_number=SATLOA2016111500118

[5] Viasat. (2018, October). *Going Global - Viasat-2 and the Viasat-3 Platform Will Take Our Service Around the World*. Available: https://www.viasat.com/news/going-global

[6] SES. (2018, October). *Exponentially More Opportunities With O3b mPOWER*. Available: https://www.ses.com/networks/o3b-mpower

[7] SES. (2018, October). *SES Selects Arianespace for Launch of SES-17*. Available:

https://www.ses.com/press-release/ses-selects-arianespace-launch-ses-17
[8] S. Finkelstein and S. H. Sanford, "Learning from corporate mistakes: The rise and fall of Iridium," *Organizational Dynamics,* vol. 29, no. 2, pp. 138-148, 2000.
[9] J. Lim, R. Klein, and J. Thatcher, "Good technology, bad management: A case study of the satellite phone industry," *Journal of Information Technology Management,* vol. 16, no. 2, pp. 48-55, 2005.
[10] E. W. Ashford, "Non-Geo systems—where have all the satellites gone?," *Acta Astronautica,* vol. 55, no. 3-9, pp. 649-657, 2004.
[11] M. Stanley, "Space: Investment Implications of the Final Frontier," 2017, Available: https://fa.morganstanley.com/griffithwheelwrightgroup/mediahandler/media/106686/Space_%20Investment%20Implications%20of%20the%20Final%20Frontier.pdf.
[12] S. Networks, "Unleashing the Potential of an Empowered - World with the Launch of O3b mPOWER," 2017, Available: https://www.ses.com/sites/default/files/2017-09/170908_SES%20Launches%20O3b%20mPOWER_FINAL.pdf.
[13] C. Balty and J.-D. Gayrard, "Flexible Satellites: A New Challenge for the Communication Satellite Industry," presented at the 25th AIAA International Communications Satellite Systems Conference (organized by APSCC), 2007. Available: https://doi.org/10.2514/6.2007-3256
[14] P. Angeletti, R. De Gaudenzi, and M. Lisi, "From" Bent pipes" to" software defined payloads": evolution and trends of satellite communications systems," in *26th International Communications Satellite Systems Conference (ICSSC)*, 2008.
[15] J. J. Knab, "Optimization of commercial satellite transponders and terminals," *IEEE Transactions on Aerospace and Electronic Systems,* vol. 49, no. 1, pp. 617-622, 2013.
[16] 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.
[17] A. París i Bordas, "Power and bandwidth allocation in multibeam satellite systems," Universitat Politècnica de Catalunya, 2018.
[18] J. Lei and M. A. Vazquez-Castro, "Joint power and carrier allocation for the multibeam satellite downlink with individual SINR constraints," in *Communications (ICC), 2010 IEEE International Conference on*, 2010, pp. 1-5: IEEE.
[19] 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.
[20] S. K. Sharma, S. Maleki, S. Chatzinotas, J. Grotz, J. Krause, and B. Ottersten, "Joint carrier allocation and beamforming for cognitive SatComs in Ka-band (17.3–18.1 GHz)," in *Communications (ICC), 2015 IEEE International Conference on*, 2015, pp. 873-878: IEEE.
[21] E. Lagunas, S. K. Sharma, S. Maleki, S. Chatzinotas, and B. Ottersten, "Resource allocation for cognitive satellite communications with incumbent terrestrial networks," *IEEE Transactions on Cognitive Communications and Networking,* vol. 1, no. 3, pp. 305-317, 2015.
[22] M. S. Net, I. del Portillo, B. Cameron, and E. Crawley, "Assessing the impact of real-time communication services on the space network ground segment," in *Aerospace Conference, 2016 IEEE*, 2016, pp. 1-13: IEEE.
[23] M. Sanchez Net, "Support of latency-sensitive space exploration applications in future space communication systems," Massachusetts Institute of Technology, 2017.

BIOGRAPHY

Markus Guerster *is a graduate student pursuing his PhD degree in the System Architecture Lab of the Department of Aeronautics and Astronautics at MIT. His interests include the development of tools to support the decision-making process for complex engineering systems under uncertainty. 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. He worked as a System Engineer at OHB System AG for 2 years besides his studies and is conducting a scientific study with in-tech GmbH to transfer aerospace methodologies to the railway industry.*

Juan José Garau Luis *is a graduate student in the AeroAstro department at MIT, currently pursuing a MSc degree. In 2017, under the CFIS program, he received two BS degrees in Telecommunications Engineering and Industrial Engineering from Universitat Politècnica de Catalunya, in Spain. Earlier that year, he was a visiting student at the System Architecture Lab, where he carried out his Bachelor Thesis on the architecture of an IoT network. He previously worked as a data scientist at Arcvi and at Barcelona Supercomputing Center. His research interests include intelligent and autonomous systems and communication networks.*

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 Lab 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 currently in orbit. Dr. Cameron received his undergraduate degree from the University of Toronto, and graduate degrees from MIT.*