Cost Analysis In A Multi-Mission Operations Environment

M. Newhouse[1] and L. Felton[2]
CSC, Marshall Space Flight Center, Alabama, 35812, USA

N. Bornas[3], D. Botts[4], and K. Roth[5]
COLSA Corporation, Marshall Space Flight Center, Alabama, 35812, USA

and

G. Ijames[6] and P. Montgomery[7]
National Aeronautics and Space Administration, George, Marshall Space Flight Center, Alabama, 35812, USA

Spacecraft control centers have evolved from dedicated, single-mission or single mission-type support to multi-mission, service-oriented support for operating a variety of mission types. At the same time, available money for projects is shrinking and competition for new missions is increasing. These factors drive the need for an accurate and flexible model to support estimating service costs for new or extended missions; the cost model in turn drives the need for an accurate and efficient approach to service cost analysis. The National Aeronautics and Space Administration (NASA) Huntsville Operations Support Center (HOSC) at Marshall Space Flight Center (MSFC) provides operations services to a variety of customers around the world. HOSC customers range from launch vehicle test flights; to International Space Station (ISS) payloads; to small, short duration missions; and has included long duration flagship missions. The HOSC recently completed a detailed analysis of service costs as part of the development of a complete service cost model. The cost analysis process required the team to address a number of issues. One of the primary issues involves the difficulty of reverse engineering individual mission costs in a highly efficient multi-mission environment, along with a related issue of the value of detailed metrics or data to the cost model versus the cost of obtaining accurate data. Another concern is the difficulty of balancing costs between missions of different types and size and extrapolating costs to different mission types. The cost analysis also had to address issues relating to providing shared, cloud-like services in a government environment, and then assigning an uncertainty or risk factor to cost estimates that are based on current technology, but will be executed using future technology. Finally the cost analysis needed to consider how to validate the resulting cost models taking into account the non-homogeneous nature of the available cost data and the decreasing flight rate. This paper presents the issues encountered during the HOSC cost analysis process, and the associated lessons learned. These lessons can be used when planning for a new multi-mission operations center or in the transformation from a dedicated control center to multi-center operations, as an aid in defining processes that support future cost analysis and estimation. The lessons can also be used by mature service-oriented, multi-mission control centers to streamline or refine their cost analysis process.

---

[1] Principal Lead Systems Engineer, MSFC/EO50, AIAA Senior Member.
[2] Systems Engineer, MSFC/EO50.
[3] Customer Service Engineering Manager, MSFC/EO50.
[4] Customer Service Engineer, MSFC/EO50.
[5] Customer Service Engineer, MSFC/EO50.
[6] Ground Systems Engineer, MSFC/EO50.
[7] Computer Engineer, MSFC/EO50.

I. Introduction

Spacecraft control centers have evolved from dedicated, single-mission or single mission-type support to multi-mission, service-oriented support for operating a variety of mission types. At the same time, available money for projects is shrinking and competition for new missions is increasing. More missions are looking for alternative low cost approaches to mission operations, and require a rapid response to inquiries about the cost of using control center services. These factors drive the need for an accurate and flexible model to support estimating service costs for new or extended missions; the cost model in turn drives the need for an accurate and efficient approach to service cost analysis.

The National Aeronautics and Space Administration (NASA) Huntsville Operations Support Center (HOSC) at Marshall Space Flight Center (MSFC) recently completed a detailed analysis of service costs as part of the development of a complete service cost model. The cost analysis process required the team to address a number of issues with broad applicability across space missions:

1) Reverse engineering individual mission costs in a multi-mission environment
2) Extrapolating costs to different mission types
3) Cost estimation for virtual, shared resources
4) Restrictions specific to government facilities
5) Cost model validation

The following sections will discuss each of these issues, relevant trades, and the lessons learned from developing and using the cost model.

Background

The HOSC has been providing operations services to a variety of customers around the world since 1960 (see Fig. 1). Its customers have included NASA flagship missions like the Hubble Space Telescope and the Chandra X-ray Observatory, small science missions like the Fast Affordable Science and Technology Satellite (FASTSAT), and the Space Transportation System (STS). The HOSC currently provides the Payload Operations Center (POC) for coordinating science operations for International Space Station (ISS) payloads, and will provide the Engineering Support Center (ESC) for the new NASA heavy lift Space Launch System (SLS).
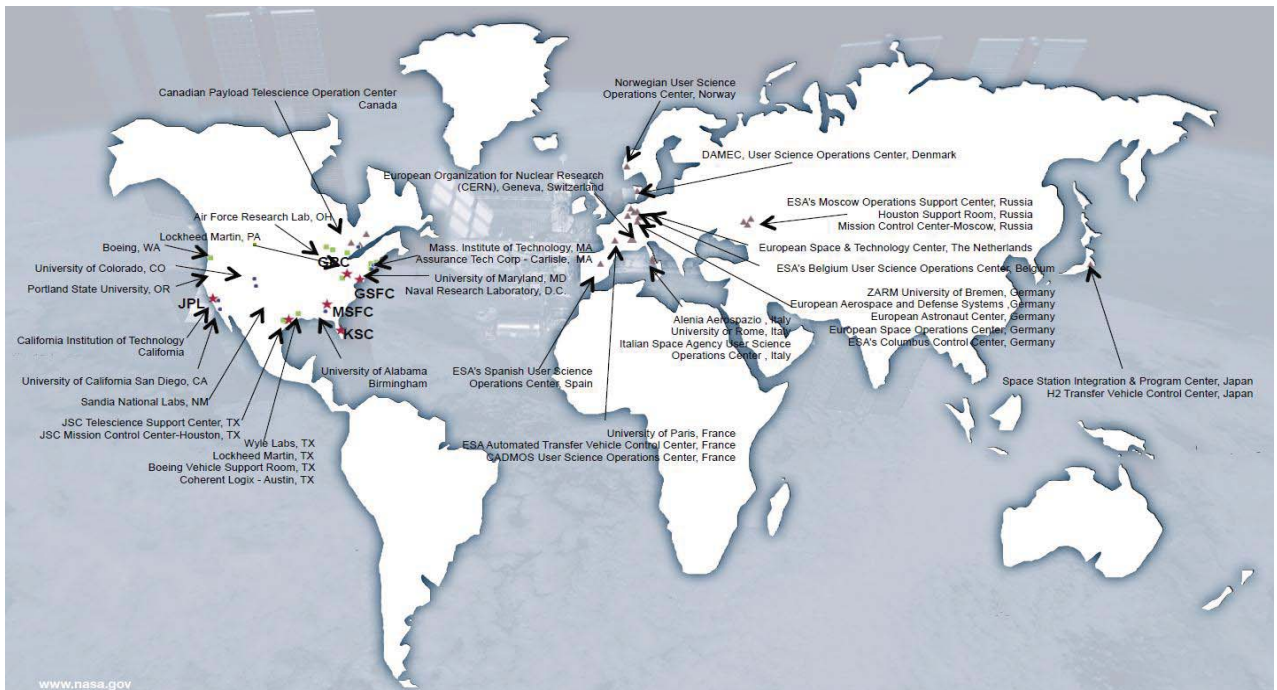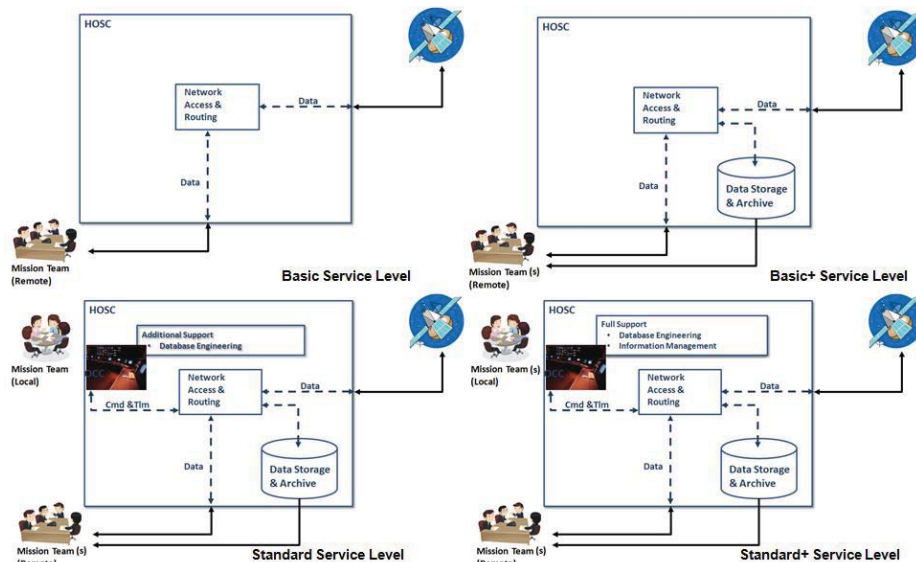

Figure 1. HOSC Customers. *The HOSC currently supports both local operations and remote operations for customers around the world.*

Given the range of mission types supported, the HOSC services have been designed to be highly configurable. Services include:

1) Secure networks
2) Data storage, retrieval, and archive
3) Database engineering
4) Command and telemetry
5) Information management

Services may support remote or local operations, and include lights out, 8x5, and 24x7 operations. Depending on the mission concept and the level of support, the services include engineering and operations labor, shared or dedicated hardware, and may include facilities. Fig. 2

Figure 2. HOSC Service Configurations. *The HOSC services have been adapted to a wide range of mission sizes and mission concepts, from a simple bent pipe for secure transmission of data between users and a space or ground asset, to full mission operations and mission data support for both local and remote users.*

illustrates the increasing support provided for four HOSC service levels; the actual services can be combined as necessary to meet the mission requirements.

The cost model is designed to provide estimates of all required resources and associated costs for the HOSC services over the full mission life cycle, including labor, hardware (including furnishings, licenses, and hardware refresh), and facilities. It starts with a high-level definition of the mission concept, mission characteristics (e.g, payload versus full mission, orbiter versus deep space), mission complexity (e.g., number of instruments, operations team size, number of operations sites), and a service level. The service levels provide the initial specification of the operational support requirements:

1) Basic Service Level includes network services only
2) Basic+ Service Level includes network and data storage, retrieval, and archive services
3) Standard Service Level includes all HOSC services, except for information management
4) Standard+ Service Level includes all HOSC services

While the service levels provide a starting point, they can be modified to refine (add, delete, or modify) the actual services and support required for a mission. The model then expands this high-level definition to the detailed facility, hardware or labor resources and associated costs. In addition, the cost model can be used to look at resources across multiple potential missions to identify conflicts in the use of resources (labor, hardware, or facilities) as new customers are added for HOSC services. Obviously, a cost model for these services must be equally configurable in order to provide accurate and competitive estimates of service costs over a full mission life cycle for a variety of missions or payloads. Specifying such a model and the algorithms supporting the model is relatively straightforward. The issues arise when populating the model parameters across the range of potential mission type and validating the model results.

II. Multi-Mission Cost Modeling Issues

A. Reverse Engineering Mission Costs in a Multi-Mission Environment

The primary customers for HOSC operations have been the ISS science payloads. Despite the distribution world-wide of remote users and the inherent uniqueness of each payload, for most there is little variation in the support provided by the HOSC for each payload. Each increment (set of 20 to 30 new or continuing payloads) undergoes the same end-to-end process of initiating accounts, identifying unique payload characteristics (if any), training new investigators, and supporting 24x7 operations for all active onboard payloads. Because of this similarity, the processes for providing services to payload customers for each new increment without interrupting ongoing payload operations have been highly optimized over the years. Hardware and labor resources are shared by all payload

American Institute of Aeronautics and Astronautics

customers, and paid for by one funding source: the ISS Program, making it difficult to separate out the costs associated with a single payload.

Support for the STS was a very different paradigm. There was one customer: the engineering team at MSFC. Operations were intermittent, driven by real-time support during each shuttle flight and access to archived data for analysis of engine performance between flights. The STS support is in the process of being replaced and re-engineered to provide similar Engineering Support Facility (ESF) capabilities for the Space Launch System (SLS). As a dedicated, single user mission, hardware and labor resources are easily isolated for the STS, but are not directly applicable to the SLS due to modernization of the technology and service approach.

The HOSC has supported several robotic satellite missions over the years, each unique. The Chandra X-Ray Observatory, one of the Great Observatory missions, was one of the first major robotic satellites operated using HOSC software, and it has been in operation since 1999. However, it is operated standalone out of Cambridge, MA, and other than the original software does not currently use HOSC facilities or services. At the other end of the spectrum, NASA's recent FASTSAT mission, was developed at MSFC in only 14 months and operated using HOSC facilities and services for its 2 year mission life. It is a prime example of rapid engineering and deployment for small satellites, but only provides one cost data point.

One of the first issues encountered while developing the cost models was reverse engineering these individual mission costs to provide meaningful information about labor resources required across the full spectrum of potential mission types. Labor data was available by customer, but primarily it was aggregate data across the customer. The most complete set of available labor data was for ISS support; however, labor data in the highly-optimized ISS paradigm was not directly applicable to other mission scenarios. In addition, mission funding (and staffing) often was capped and therefore support was resource-limited, rather than being product-driven, and the goal was to avoid propagating these limitations to future missions. The available labor data was highly variable over time, with the high-level of support around key events (gate reviews or major technical interchanges), and significantly less support between these events, making extrapolation to different mission development schedules difficult. To compensate for these issues, we used a combination of analysis of actual data for past and current missions, and detailed interviews of engineering leads concerning the activities and products for supporting missions of various types throughout each phase of the mission life cycle. The actual labor data included engineering and operations support and product development; the detailed interviews covered level-of-effort support to mission systems engineering, product development, operations support, and sustaining engineering. In general as shown in Fig. 3, the actual labor costs were found to provide a lower limit on the resources required for a mission; the detailed interviews provided an upper limit.

Figure 3. Staffing Level Analysis. *Sample plot showing the actual (aqua through light green lines) staffing support for a given engineering position versus the calculated support for that position based on detailed interviews (blue line). Expected staffing is bounded by the actual and estimated values.*

From another perspective, performance analysis tools can be used to provide data on the peak and average loading on hardware and software systems. The resource utilization for a given customer, for example an ISS payload, was easily obtained using these tools. Ideally, it would be desirable to be able to analyze this data by user and user role, to assist in extrapolating the resource loading to larger and smaller, or more or less complex missions. In reality, breaking the data down to the individual user level in order to extrapolate to missions or payloads of different size or complexity was more difficult and time consuming. Instead the data needed to be looked at for "classes" of missions or payloads with similar characteristics. In addition, performance tools only provide data on performance for the system of today; they provide no hard data on the performance of the system of tomorrow. This must be extrapolated based on expected improvements in technology and, just as critically, changes in user operations concepts and expectations.

The final model relied heavily on the expertise of the customer service team (CST), which supports a HOSC customer from day one through mission closeout, to understand the data and make the appropriate comparisons between the available cost data and the mission type, size, support requirements, and complexity.

American Institute of Aeronautics and Astronautics

B. Benefits and Risks of Modeling in a "Cloud-like" Environment

"Cloud" is the word of the moment. Ideally, the most efficient use of resources (both hardware and labor) can be made when resources are shared in a virtual environment. Rather than requiring each mission to purchase dedicated hardware, hardware that typically must be redundant to meet reliability and availability requirements, sharing hardware reduces the cost commitment of each mission, particularly smaller missions, without increasing the risk to mission performance. A virtual hardware environment also mitigates the risks associated with inaccurate estimates of the resources required for each customer, or variations in the resource requirements over the life of the mission. Slight over-estimates of the resources required for a more complex mission can be offset by slight under-estimates of resources required for simpler missions. Likewise, virtual sharing of resources means that any individual mission does not have to scope resources for peak loading. Instead, resources can be appropriately allocated based on mission phase and critical activities. Of course, this is limited by and must allow for overlapping peak support when critical activities align for multiple missions.

However, commercial cloud providers have an advantage in that they can amortize the cost of maintenance and refresh on an annual or monthly basis over all customers and bank the money until it is required. The money is then spent on shared equipment used by all customers. Government installations do not have this advantage. Funding is typically level and cannot be held until needed. Funds must be spent within a specific timeframe, typically within a year or two of allocation. Funding is provided for each customer, and must be allocated and spent in support of that customer.

The discrete nature of hardware adds additional complexity to the problem. Server or network capacity can be increased when support for a new mission will exceed capacity. But the cost may not be proportional to the mission size. For example, as illustrated in Fig. 4, currently funded missions may be using most or all of the existing system capacity; a small mission may only require a small percent of the overall capacity, but this is sufficient to require the purchase of new equipment. The cost of this new hardware cannot be shared with the current mission, which is working to existing, approved funding levels. However, by the very nature of its low cost profile, the small mission cannot afford to purchase an appropriate cloud-like increment in capacity. In this case, purchasing dedicated equipment for the smaller mission may be the most cost effective approach, even if it means the mission cannot take advantage of the economies of scale.

Figure 4. Hardware Resource Example. *Resource requirement for a small mission may exceed existing capacity, but increasing the resource incrementally may exceed the small missions budget for operations support.*

There is no clean solution to this dilemma. The HOSC cost model takes a hybrid approach that can take advantage of shared resources when mission concept, capacity, and cost permit, but can also fall back to the traditional case of dedicated, low-cost hardware if necessary.

C. Cost Model Validation and Extrapolation

The HOSC has supported and continues to support a variety of missions. Ideally, validating the model across the potential range of missions would involve comparing estimated resources and associated costs against truth models for each type of mission. In actuality, detailed actual labor costs and hardware resource usage are only available for the more recent customers and missions. In addition, the hardware resource metrics are geared toward ensuring service levels and availability requirements are met in a multi-mission environment, as well as early identification of potential issues, rather than toward the more difficult mapping of resource usage to individual missions, users, or user roles. The second choice would be validating the estimates against actual data collected for new missions, and updating the model in an Agile-like approach. In practice, the long development cycle for most NASA missions makes this approach difficult. Add to this the decrease in the number of funded NASA missions, and this approach becomes impractical.

However, in the tight NASA budget environment, estimated costs need to represent a realistic, not-to-exceed limit on the cost of designing, testing, and executing mission operations. As in other areas, we have taken a hybrid approach to validation. We have modeled current or completed missions with the accuracy of the available data. New estimates, and the mission assumptions that underlie the estimate, are subject to review by cognizant engineers to ensure consistency with past experience as well as ongoing or expected architecture and service enhancements.

More importantly, we are in the process of reviewing and revising the measurements collected and the metrics used to analyze operations performance and capability. The metrics need to clearly map resource use (labor and hardware):

    1) Mission characteristics
    2) Mission phase
    3) Critical activities
    4) Users and user roles

Since its completion, the model has been used to cost operations support for missions that include a lunar surface operations, low earth orbit operations for cube satellites, deep space operations for asteroid rendezvous, and a backup control center for an in-flight mission. To date, those estimates have demonstrated a high-level of agreement with independent estimates by the system engineers and engineering management.

III. Conclusion

The variety of missions supported by the HOSC continues to evolve. The tight NASA budget requires accurate, realistic operations cost estimates covering the life of the mission, while the realities of competition for missions requires a rapid, low-cost response. The HOSC cost model meets these needs. However, developing the model has identified a number of lessons learned that can be applied to developing multi-mission cost models, or cost models in general. Note that while this paper specifically addresses cost analysis in a multi-mission environment, these lessons apply equally to extrapolating costs for a new mission based on operation of an existing mission, in a single-mission environment.

First, the customer of tomorrow is not necessarily the customer of today. When setting up metrics for an operational system, the use of those metrics needs to be considered. The simplest measurements derived from performance analysis tools may be adequate for monitoring performance against customer requirements, but not provide adequate granularity in a multi-mission environment to allow the data to be extrapolated to other missions. Additional measurements, or the ability to correlate the measurements to other characteristics, will be needed to truly characterize the resource requirements across a spectrum of mission types.

Second, multiple approaches and the ability to synthesize data across those approaches (for example, interviews and data modeling), are necessary to extrapolate from a limited set of data points (individual missions with specific characteristics) to a broad spectrum of mission types.

Third, changes in technology and operational use of the system need to be considered and incorporated in the model parameters or design.

Fourth and foremost, the involvement of an experienced and knowledgeable engineering team is critical to understanding the data and its applicability in an evolving service environment. Cost analysis needs to be tempered by the detailed understanding of user needs, engineering reality, and the ongoing evolution to the operations center of tomorrow.

Acknowledgments