Development of the ITACA Network Loading Analysis Tool's Scheduling Techniques

Joseph M. Oglio and Bryan W. Welch
Glenn Research Center, Cleveland, Ohio

NASA STI Program . . . in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA Scientific and Technical Information (STI) Program plays a key part in helping NASA maintain this important role.

The NASA STI Program operates under the auspices of the Agency Chief Information Officer. It collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI Program provides access to the NASA Technical Report Server—Registered (NTRS Reg) and NASA Technical Report Server—Public (NTRS) thus providing one of the largest collections of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counter-part of peer-reviewed formal professional papers, but has less stringent limitations on manuscript length and extent of graphic presentations.

- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., "quick-release" reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.

- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.

- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.

- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

For more information about the NASA STI program, see the following:

- Access the NASA STI program home page at http://www.sti.nasa.gov

- E-mail your question to help@sti.nasa.gov

- Fax your question to the NASA STI Information Desk at 757-864-6500

- Telephone the NASA STI Information Desk at 757-864-9658

- Write to:
  NASA STI Program
  Mail Stop 148
  NASA Langley Research Center
  Hampton, VA 23681-2199

Development of the ITACA Network Loading Analysis Tool's Scheduling Techniques

Joseph M. Oglio and Bryan W. Welch
Glenn Research Center, Cleveland, Ohio

Acknowledgments

Level of Review: This material has been technically reviewed by technical management.

Development of the ITACA Network Loading Analysis Tool's Scheduling Techniques

Joseph M. Oglio* and Bryan W. Welch
National Aeronautics and Space Administration
Glenn Research Center
Cleveland, Ohio 44135

Abstract

NASA's SCENIC project aims to simplify and reduce the cost of space mission planning by creating analysis capabilities which are integrated with relevant analysis parameters specific to SCaN assets and SCaN supported user missions. The Integrated Tradespace Analysis of Communications Architectures (ITACA) will provide an all-in-one package for various analysis capabilities that normally require add-ons or multiple tools to complete. The ITACA tool will be responsible for assessing the given network architecture and generating a schedule for the missions as well as the assets. ITACA will allow users to evaluate the quality of service of a given network and determine whether or not the network will satisfy the mission's requirements. ITACA is currently under development, and during the spring of 2018 major development in the tools scheduling techniques were completed. A total of seven different techniques were completed.

Nomenclature

GRC   Glenn Research Center
GSFC   Goddard Space Flight Center
ITACA   Integrated Tradespace Analysis of Communications Architectures
NASA   National Aeronautics and Space Administration
SCENIC SCaN Center for Engineering, Networking, Integration, and Communications
SCaN   Space Communications and Navigation
UI   user interface

Introduction

The goal of NASA's SCENIC project is to provide users of NASA's SCaN networks (a given collection of assets used to service a mission) with a free web-based user interface (UI) to perform on-the-spot mission analyses. SCENIC, in providing a comprehensive UI with advanced analysis capabilities, will save its users both time and money in the early stages of mission planning using NASA's SCaN networks. One of the capabilities SCENIC wants to provide to its users is the ability to evaluate the service of existing and custom SCaN networks for mission satisfaction. When analyzing the capabilities of a network four metrics are assessed: Mission Data Volume, Mission Latency, Network Capacity, and Network Utilization. These metrics will vary with the number of missions being provided service and the requirements of those missions. Network analysis is important for creating a schedule that will satisfy the requirements of as many missions as possible given the requirements of the missions and the specifications of the network assets. The tool SCENIC will use to perform loading analysis and create the network service schedule is known as the Integrated Tradespace Analysis of Communications Architectures (ITACA).

The objective of the internship was to develop and implement new scheduling techniques to the MATLAB®-based (The MathWorks, Inc.) ITACA network loading analysis tool. By factoring in sets of mission requirements and network asset specifications during the timeframe of interest, ITACA will create an optimized schedule using real-world operational parameters, as well as provide users with information detailing which network assets are available to service their missions, and when those assets will be able to establish and end contact (contact window). Additionally, ITACA will provide users with assessment metrics—effectively a quantitative scorecard that allows users to assess the ability of a given SCaN network architecture to provide proper mission service and satisfy the mission requirements. Throughout the course of the internship, various scheduling techniques were created that are geared to optimize one or many of the metrics assessed in the tool.

Methods and Results

Creating a scheduling technique for the ITACA tool that is fast, accurate, and optimized is a difficult task. Any technique that is created is expected to analyze a network of over 100 assets and 100 missions, which can easily generate a list of over 10,000 total access windows. With this vast amount of windows the number of possible schedules exceeds 10000 factorial which is over 2.8e35657. The ultimate goal of these scheduling techniques is then to produce the best possible schedule from the 2.8e35657 options. When creating an algorithm to do there are two things that need to be considered, the first is time, and the second is the results. The amount of time spent creating the analysis can be determined by the method you choose as can the results. It isn't difficult to create a working schedule, but to create the best possible schedule from the windows given can take a lot of work. If one window is chosen, then the time used by that mission as well as that asset can no longer be used in any future piece of the schedule, assuming the asset is not multiaccess (can communicate with multiple missions) and that the windows given do not include any repeaters (nodes in the path of communication other than the start and end node). So, all the remaining windows must be cut to accommodate the scheduled ones in order to ensure no overlapping in the schedule. Picking one window therefore effects the remaining windows making the schedule interdependent with what has been previously scheduled. This makes it nearly impossible to predict what the schedule will look like in the end, which further leads to it being nearly impossible to predict the final scores of this schedule. The scores are the important part of the result that we want to make as high as possible. The only way to know for sure that you have the highest scores possible would be to create every other possible schedule, generate the scores of all of them, and finally, compare to find the one with the highest score. This process would take an unrealistic amount of time, so the best process of creating the schedule must balance the amount of time it will take, and allowable score variability.

The scheduling problem is a more complicated version of the traveling salesman problem. In the traveling salesman problem, a salesman is tasked with visiting every city given taking the shortest path possible, starting and ending in the same city, and visiting every city only once (other than the starting/ ending city). Other variations of this problem involve changing the goal from taking the shortest path possible by distance, to time, providing the times it would take to go from one city to another. In the scheduling problem however, there are many salesman (missions), going to many cities (assets), with travel times (latency), sales rates (data rates), the amount of time spent in each city (duration of the link), and unless a city is big enough (multiaccess) only one salesman can be in a city at a time. With these problems being very similar, techniques used in the traveling salesman problem can be adjusted to fit the scheduling problem.

The original scheduling technique and the alternate techniques created during this internship are all variations of greedy algorithms. This type of algorithm looks at a subset of data and picks the current best from those options in an effort to reach the goal. A greedy algorithm is one of thousands of techniques that could be used to solve this problem. The reason that this approach is the one used is because it produces a schedule close to the solution. In the traveling salesman problem a greedy approach picking the nearest neighboring city as the next destination results in a path 25 percent longer than the shortest path (Ref. 1). This leads to a belief that a good greedy algorithm could lead to a schedule that is 25 percent worse than the best possible schedule on average. Of course 25 percent is only an average, given specific data points it can be shown that the same greedy algorithm can also produce the worst possible path, or schedule (Ref. 2).

In the algorithms designed the subset is varied and the overall goal is varied, producing many very different schedules. Picking specific goals can also make the greedy algorithms closer than 25 percent away from the best solution. One factor that contributes to the scores is sometimes mission priority. Certain mission's needs and goals are more important than others, so even if the schedule produced isn't the best global schedule it might be the best schedule for the top mission priority levels.

The Data Volume Satisfaction score is created by comparing the data volume achieved for a mission to the data volume requirement for that mission. If the data volume achieved is higher than or equal to the requirement than that mission gets a 1 for satisfied, if not the mission is given a 0 for not satisfied. The percent of 1's is then the Data Volume Satisfaction score.

The Latency Satisfaction score is based on achieved latency and latency requirement. The latency is the amount of time the mission goes without contact. If the latency is higher than the latency requirement than that mission failed at meeting its goal and is given a 0. Like the Data Volume Satisfaction score the matrix of 1's and 0's is used to determine the percent of goals met.

The Network Capacity score is determined by the percentage of the maximum bandwidth of the assets used in the schedule (bandwidth utilized). The bandwidth utilized is determined by the modulation, coding scheme, and data rate of a given link in the schedule. The average of all of these percentages produces the overall Network Capacity score, which shows if the assets in network are being used effectively.

The Network Utilization score is a two part score. The first piece is the percentage of the interval that the assets given are used according to the schedule. The more they are used the higher the percentage is. With these four metrics and multiple scheduling techniques one can determine if the network is too weak to meet the needs of the missions. With multiple greedy scheduling techniques there should not often be cases where the schedule produced is far from the best possible schedule.

During the internship session several scheduling techniques were created geared to meet these four metrics as best as possible given the list of access windows and the requirements. As discussed earlier all of these techniques are greedy algorithms meaning they do not optimize globally, but instead try their best to optimize one or only a few goals. The first four techniques are all similar in that they choose a subset of windows to analyze and then pick the window with the highest data volume. The windows shown are color coated to indicate their data rate as shown in Figure 1.
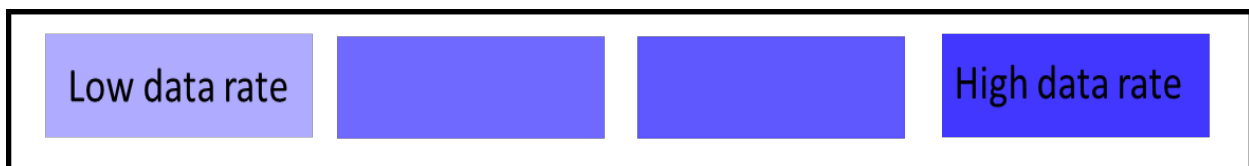

Figure 1.—This image displays information about the data rate of the access windows used in the future images. As the image suggests the darker the purple the higher the data rate will be during the event.

Data Volume Approach 1

The first of these techniques created was a primitive one. This technique is entirely priority based and chooses windows from the first priority mission to the first priority asset. Once this combination is completed it moves on to the next asset priority level, and once those have been exhausted it moves to the next mission. Since this priority based approach only looks at windows from specific mission and asset combinations as seen in Figure 2 it is unlikely to produce desirable scores. However, if priority is the top priority, this technique will deliver.

Data Volume Approach 2

The next approach designed was to be more global than the previous in that it only prioritized against the assets. This approach was designed to improve upon the Network Utilization score by only using the windows with high data volume as seen in Figure 3.
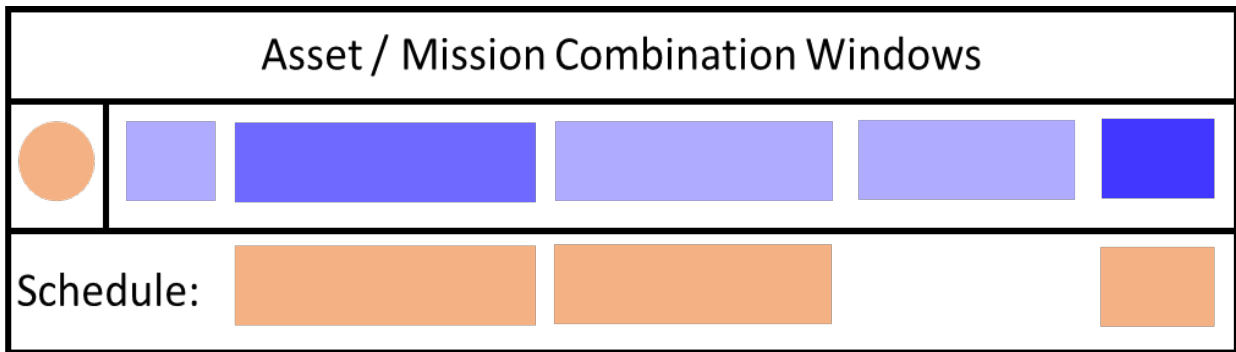

Figure 2.—This image displays an example of what the schedule would look like for a single mission and asset combination using the priority based scheduling technique.
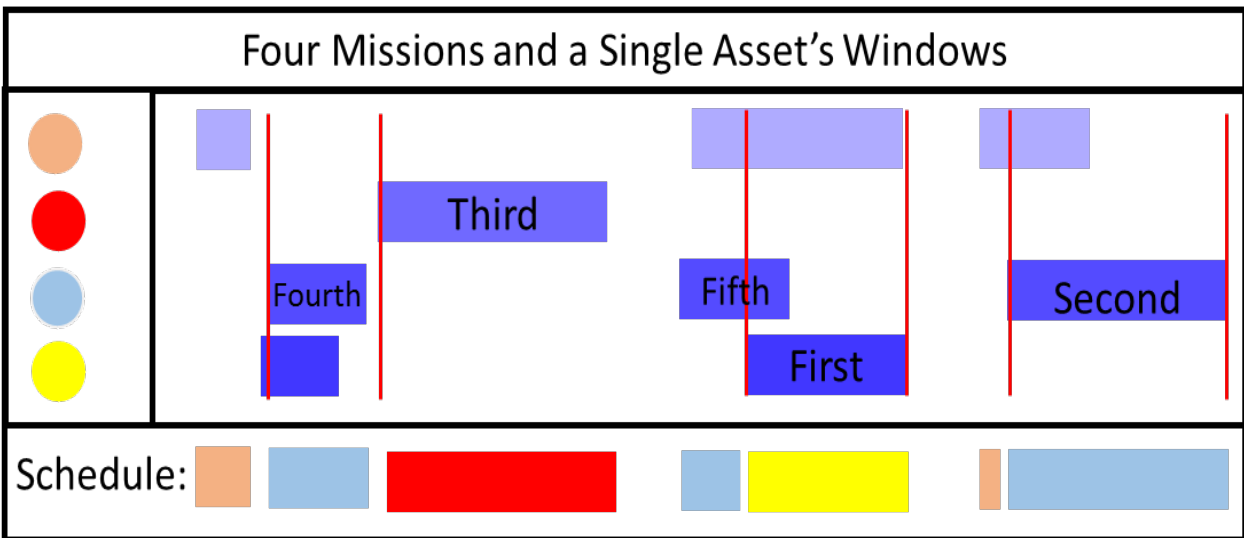

Figure 3.—This image displays an example of what the schedule would look like for a single asset when scheduled to support four missions using the asset priority approach. The major windows are numbered in the order they are picked as they will cut the other access windows as shown by the red lines.

Data Volume Approach 3

The third technique focuses on mission priority rather than asset priority. Which means that the windows analyzed are all of the windows used by a specific mission. If the higher priority missions have enough contact times to meet their Data Volume Requirements this technique ensures they will. This approach will hurt the lower priority missions that will be scheduled after the higher priority missions take the high data rate windows. It also is likely to produce large latencies when the mission reaches its Data Volume Requirement as shown in Figure 4 by the yellow circle.

Data Volume Approach 4

The fourth scheduling approach was designed to increase the total data volume sent across the network. Instead of going in order of priority this approach analyzes every contact window at the same time and chooses the one with the highest data volume. This approach as shown in Figure 5 meets the Data Volume Requirement of the top mission. It also appears to have satisfactory marks in the other scores, for the first mission. For the second mission however, things are a bit different. A huge latency is present near the end of the contact as shown by the blue circle. The mission also uses very few contact windows which makes it likely to come short of the Data Volume Requirement.
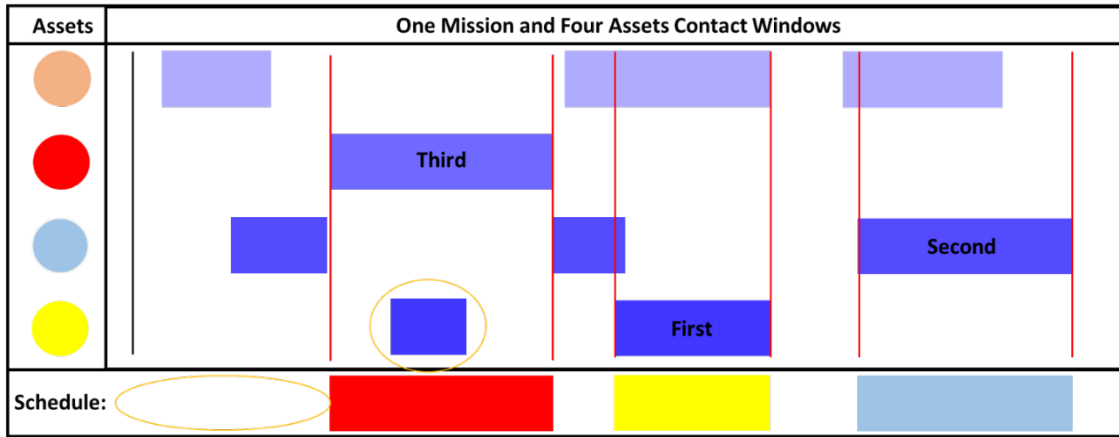

Figure 4.—This image displays an example of what the schedule would look like for a single mission using four assets and the mission priority scheduling approach.
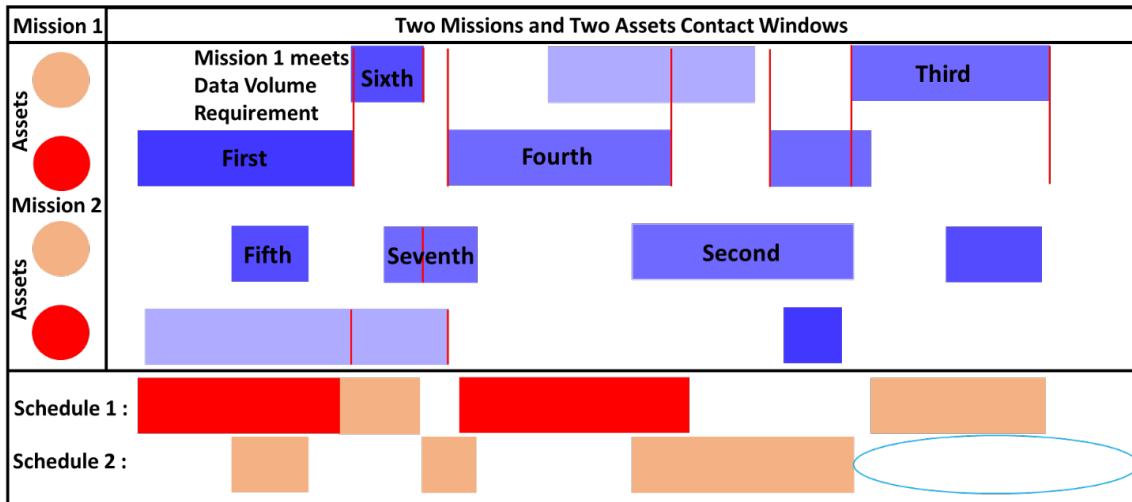

Figure 5.—This image displays an example of what the schedule would look like for two assets and two missions using the global scheduling approach.
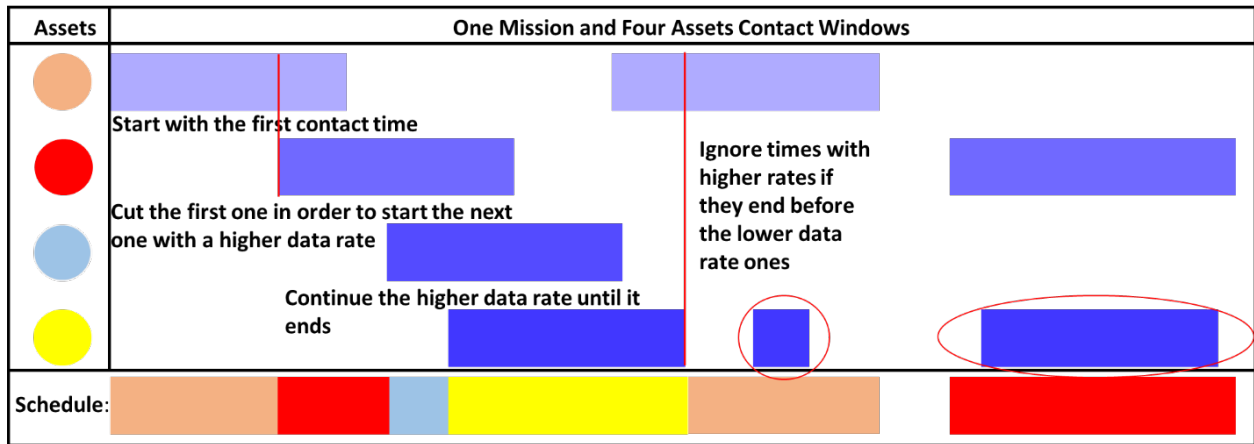
Figure 6.—This image displays an example of what the schedule would look like for one mission and four assets using the latency scheduling approach. The red circles indicate possible optimization issues for the technique.

Latency Approach

The fifth approach chooses the subset of windows by mission priority like the third method, but this approach no longer aims directly for the Data Volume Satisfaction score. Instead it is designed to increase the Latency Satisfaction score. This approach as shown in Figure 6 uses more complicated logic when determining which window to pick from this subset. This approach though does have a potential downside, as seen in the circled windows of Figure 6. These windows have very high data rates but are not chosen because they wouldn't help the latency even though they do help the data volume.

Realistic Approach

The next approach uses the same technique as the previous scheduling approach. The difference between the two is that after the selected windows are chosen and integrated into the schedule the first approach cuts the remaining windows to ensure none overlap. This approach does not do that. Instead the full list of windows are used to create a request for that mission. Then in priority order these requests are compared to the schedule. If there are overlapping times they are removed and not cut. The remaining windows are then added to the schedule. These processes can be seen in Figure 7 and Figure 8. After every mission has been added to the schedule the final process of dividing up the unused time (TUT) begins. The entire list of windows is cut by the scheduled ones. Next, missions are selected in random order and use the TUT time that has the highest data volume, assuming they haven't met their Data Volume Requirement yet. This approach is not an optimal one as parts of it use randomization, but this approach is designed to simulate how scheduling is currently done. The process of requesting time in priority order and getting that time accepted or rejected is the approach currently used by the SCaN team.

Hybrid Approach

The final scheduling approach was designed to accommodate the Latency Requirement and then the Data Volume Requirement. This approach uses most complicated logic by far when compared to the other techniques, but it is also likely to produce the best scores. The description of how the approach works as well as an example of it can be found in Figure 8. The thick black lines and the red lines are used to indicate the Latency requirement of this mission. The window in the bottom right hand corner is circled because it may or may not be scheduled. The red line from the previous window is further out than the end of the interval so the Latency Requirement has been met in that area. The Data Volume Requirement however, may or may not have been met. If not then this window is utilized in the schedule as shown.
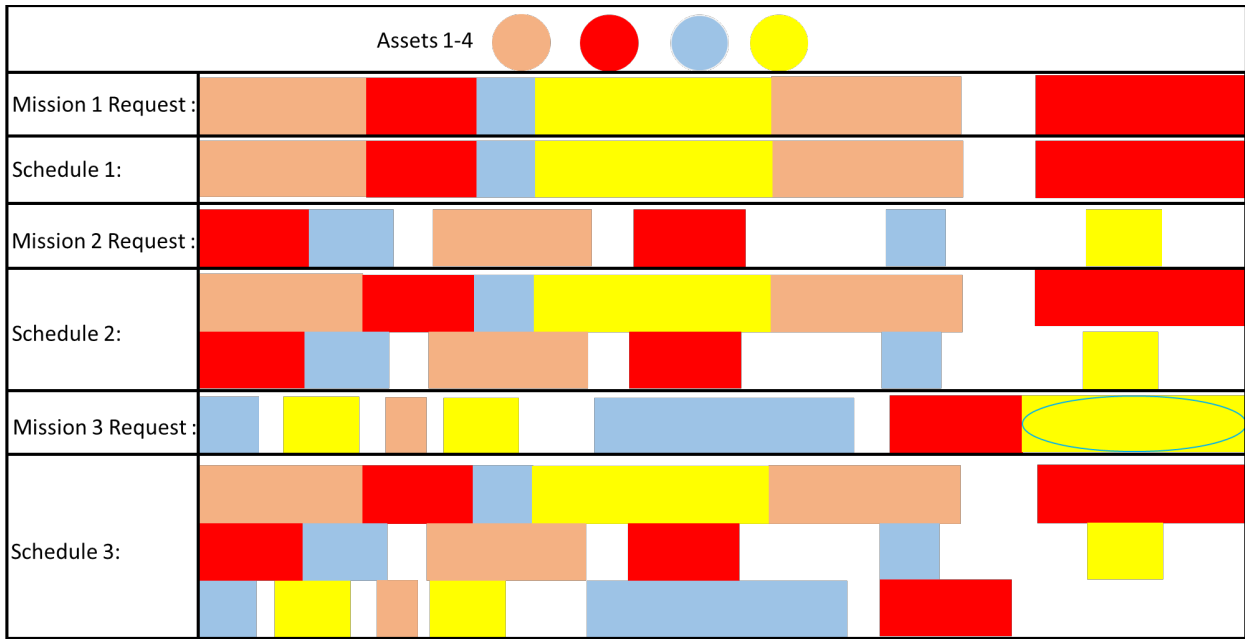
Figure 7.—This image displays an example of what the schedule would look like for three missions and four assets using the real scheduling approach.
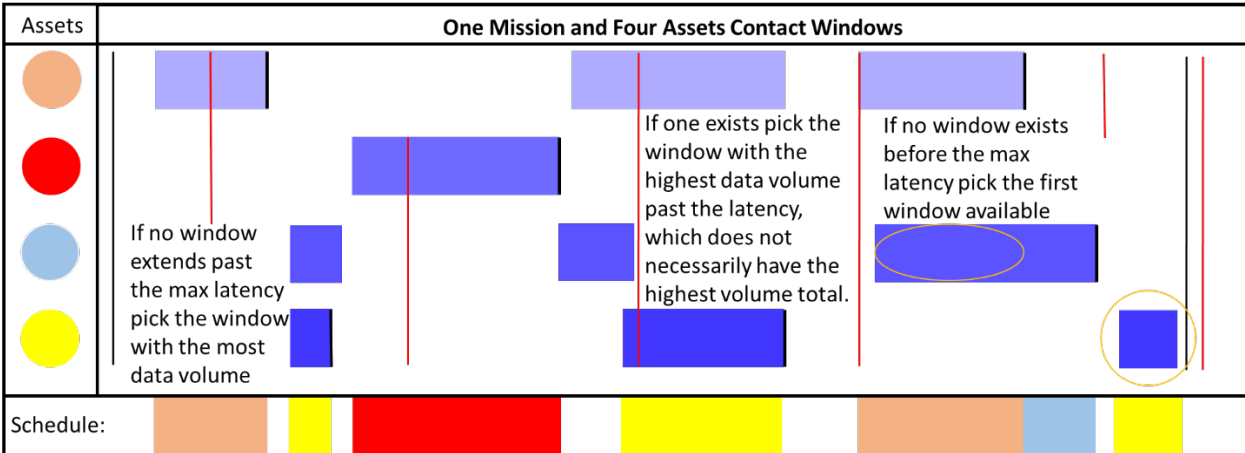

Figure 8.—This image displays an example of what the schedule would look like for one mission and four assets using the hybrid scheduling approach. The first circled window represents a possible optimization issue as it has a much larger data rate.

If none of the schedules from the alternate scheduling techniques used produce a desirable result the tool is designed to also output the list of unscheduled, uncut, windows that the schedules were created from. The user can then use a different scheduling technique if they so desire.

Conclusion

The ITACA network loading analysis tool will be used by SCENIC to assess the capabilities of the network given a defined mission set. ITACA creates both an optimized network service schedule and quantitative assessment metrics that allow the user to determine whether or not a certain architecture will satisfy the mission requirements. Both of these features will help users decide if the SCaN program has the necessary capabilities to service their mission(s) properly. Given the scope of ITACA in SCENIC, the improvement efforts of the internship will have a direct impact on the future plans and capabilities of

SCENIC. Furthermore, some of the efforts of the internship will be applied to other tools and capabilities within SCENIC. Despite the success of improvements made to ITACA over the course of the internship, further improvements must still made.

One improvement to the tool that can be performed is to integrate repeater capabilities. In order to do this changes would need to be made to other algorithms as well but especially the scheduling algorithms. In order to implement repeaters correctly the storage capabilities of the assets would need to be known and tracked across time as the schedule was created.

In addition to upgrading the scheduling techniques testing needs to be done on them. The seven designed techniques need to be run with many different analyses so the scores can be compared. The schedules produced by the realistic approach also need to be compared to the real communication schedules to determine if it accurately uses their methods.

ITACA is collaboratively developed by personnel at both NASA Glenn Research Center (GRC) and NASA Goddard Space Flight Center (GSFC), and SCENIC will continue the development of ITACA with the hope that it will be ready for integration into the UI in late 2018.

References

1. Johnson, D. S.; McGeoch, L. A. (1997).The Traveling Salesman Problem: A Case Study in Local Optimization. *In Arts, E. H. L.; Lenstra, J. K. Local Search in Combinatorial Optimization. London: John Wiley and Sons Ltd.* (pp. 215–310).
2. Gutina, Gregory; Yeob, Anders; Zverovich, Alexey (15 March 2002). Traveling salesman should not be greedy: domination analysis of greedy-type heuristics for the TSP. *Discrete Applied Mathematics*. 117: 81–86.