

# 3 Petabytes or Bust – Planning Science Observations for NISAR

Joshua R. Doubleday

Jet Propulsion Laboratory, California Institute of Technology, {firstname}.{lastname}@jpl.nasa.gov

## ABSTRACT

The National Aeronautics and Space Administration (NASA) and the Indian Space Research Organization (ISRO) have formed a joint agency mission, NASA ISRO Synthetic Aperture Radar (NISAR) to fly in the 2020 timeframe, charged with collecting Synthetic Aperture Radar data over nearly all of earth's land and ice, to advance science in ecosystems, solid-earth and cryospheric disciplines with global time-series maps of various phenomenon. Over a three-year mission span, NISAR will collect on the order of 24 Terabits of raw radar data per day.

Developing a plan to collect the data necessary for these three primary science disciplines and their sub-disciplines has been challenging in terms of overlapping geographic regions of interest, temporal requirements, competing modes of the radar instrument, and data-volume resources. One of the chief tools in building a plan of observations against these requirements has been a software tool developed at JPL, the Compressed Large-scale Scheduler Planner (CLASP).

CLASP intersects the temporo-geometric visibilities of a spaceborne instrument with campaigns of temporospatial maps of scientific interest, in an iterative squeaky-wheel optimization loop. While the overarching strategy for science observations has evolved through the formulation phases of this mission, so has the use of CLASP.

We'll show how this problem space and tool has evolved over time, as well as some of the current parameter estimates for NISAR and its overall mission plan.

KEYWORDS: observation scheduling planning modeling

## 1. BACKGROUND

### 1.1 Mission

The NISAR mission is the Synthetic Aperture Radar (SAR) portion of of NASA's answer to the National Research Council's (NRC) Decadal Survey for previously unavailable data and insight in three earth science domains: Deformation, Ecosystem Structure, and Dynamics of Ice (DESDynI) [1]. The mission concept has undergone several drastic revisions over the years such as from a two-spacecraft mission with an additional LIDAR instrument, to the current one-spacecraft design, and the collaboration with an international partner; The Indian Space Research Organization (ISRO) is partner to NASA's Jet Propulsion Laboratory (JPL) on this mission. ISRO is to provide the mission Launch services and vehicle, spacecraft bus, spacecraft operations infrastructure, and an S-band SAR instrument. JPL is providing the L-band SAR instrument, radar reflecting antenna and structure, GPS and solid-state recorder. Each agency will provide high-rate telecom systems to receive science data transmitted direct-to-earth stations.

These two SAR instruments combined will produce data rates of upwards of  $5 \times 10^9$  bits per second for intervals on order of several minutes, and sustained rates for global mapping of deformation objectives on order of  $1-2 \times 10^9$  bits per second. These projected rates and duty test the bounds of today's infrastructure for moving this data from orbit to the ground.

Mission science requirements have been categorized into campaigns, by scientific discipline, regions of interest, mode of operation of the instruments, and temporal constraints on observation frequency. There are upwards of 25 such disciplines, overlapping to varying degrees in space, time, and operational modes.

The spacecraft observatory is currently baselined for a 12 day repeat cycle polar orbit, sun-synchronous for dusk and dawn passes.

### 1.2 CLASP

To enable trade studies of the mission science objectives, and provide a basis for system design, schedules of observations are planned with CLASP[8], a software scheduling and planning tool with an emphasis on planning against geometric constraints, in addition to typical temporal and resource constraints. While many software planners allow an arbitrary hierarchy of activities, CLASP keeps a rigid structure of spacecraft and sensor definitions that result in a finite set of parameterizations of "observation" activities that are the main actors on the schedule [6]. Amongst the landscape of software planning tools described in [3], CLASP shares many common scheduling features, to varying extents. CLASP models resources as either depletable or, non-depletable, or both where the former is the integral of the former. CLASP

also models states, however these are somewhat restricted in domain rather than completely user definable, for example spacecraft attitude, or instrument mode.

This rigid planning schema trades the flexibility of a general purpose planner for more power in the domain of spatial reasoning, and applicability towards planning coverage for mapping missions. SciBox is a similar software tool from Johns Hopkins Advanced Physics Laboratory. SciBox was used on the MESSENGER mission for its geometric/geographic capabilities for planning large-area mapping [4] [5], though compared to CLASP it emphasizes a broad toolkit of generic spacecraft planning capabilities entirely written in JAVA. Both software tools feature modules to compute observation opportunities from science campaigns and constraints, as well as mechanisms to optimize the schedule of observations. Unlike SciBox, CLASP leverages several existing libraries of code and applications such as SPICE for ephemeris calculations and Google Earth for visual rendering. While SciBox is a toolkit requiring development for use, CLASP can generate useful simulations out of the box, though complex models will also require adaptation.

CLASP will optimize a schedule of observations given a set of spacecraft, spacecraft trajectories, instruments, instrument modes, mode compatibilities/dominance, datarates, data storage parameters, downlink schedules and rates, and finally sets of geometric target campaigns with desired temporal constraints including windows of opportunity and repetitions, with each target assigned a scoring weight.

Sensors are modelled as pushbroom style sensors, sweeping out coverage swaths from the moment an observation begins to end of observation. Geometry of the sensors are parameterized by near and far look angles as angles rotated about the velocity vector of the spacecraft from the nadir look vector, looking 90 degrees off of velocity. The swath can be further parameterized to point with a rotation about the nadir vector. The sensor may operate in any number of modes each constrained to states of the spacecraft, and dictating a data-generation rate on the spacecraft storage system.

The spacecraft is defined by the provided ephemeris or trajectory model, a target body defaulting to earth, and encompasses a state timeline for the sub-components. A simple monolithic solid-state recorder houses instrument data until released by playback on a downlink schedule. The spacecraft maintains a state for ascending vs descending and left vs right looking to constrain operation of modes of the instruments. The spacecraft has a parameterized ability to dynamically roll about its velocity axis adjusting the viewing geometry of the sensors. The timeline of the spacecraft roll is populated as part of the planning process.

Science campaigns are described by geographic regions of interest (ROI) in keyhole markup language (KML). Within a description field of a region of interest within the KML file CLASP parses a campaign description language. A single ROI may have any number of science campaigns. Each campaign describes a required instrument and mode, a scheduling priority, a weight to the campaign used to assess the aggregate utility of a generated plan, geometry constraints on the observation in terms of left vs right looking and ascending or descending node. Lastly a campaign will require some defined number observations within a collection of time windows.

To plan against the science campaigns, CLASP projects the sensor geometry definitions against the spacecraft trajectory onto the target body over time at discrete time intervals, sweeping out a schedule of swaths. The regions of interest are then either directly intersected with all swaths generating a collection of shards of target visibilities over time, or by first projecting the target onto a regular spacing of gridpoints on the target body. The former solution gives coverage calculation at resolution limited only by processor memory and floating-point accuracy, however for scenarios of sufficient planning horizon and/or where the swath may intersect with itself and thus forming many shards it becomes intractable and the later solution necessary.

There are a number of scheduling algorithms CLASP can utilize to generate the schedule, but its chief approach is to utilize the Squeaky Wheel Optimization algorithm, coupled with a simple greedy sweep forward in priority order of campaign targets. Solutions are iterated over some finite number of attempts to find a maximum score while adjusting target priorities internally between iterations.

## 2. APPLICATION

Throughout the preliminary design phase (phase-B) of NISAR development we have continued to refine the plan of observations to achieve the mission science requirements. In contrast to earlier phases of mission development, coarse parameters such as the reference orbit, downlink strategy and stations, solid-state-recorder (SSR) size, have begun to settle, and efforts have been focused on creating a higher-fidelity model. We will discuss how CLASP has adapted to this evolution of usage, while presenting some real results in the context of this mission.

## 2.1 Coarse/Rapid analysis for fast-repeat cycle studies

NISAR has chosen an orbit of approximately 747 km altitude and a slightly retrograde orbit to provide an exact repeat ground-track every 12 days as its baseline reference orbit. However, this represents a compromise to some science objectives such as studying fast-moving ice which would prefer to observe much more frequently and the challenge of observing all land-masses with a finite field of view (requires a longer minimum temporal baseline), as well as presents a pressure on the observatory's initial checkout timeline after reaching orbit as any single calibration site in the non-polar latitudes are generally only observable 2 times (from two different geometries) every 12 days. As such, a study had been conducted to study the feasibility of temporarily leveraging a different orbit to accelerate the repeat period from 12 days to 5 or 2 days.

CLASP was able to support this study in terms of defining coverage of science and/or calibration targets completely off the shelf since in this case we were not interested in demands of resources specific to NISAR, e.g. the NISAR SSR, and thus not necessitating the NISAR adaptation of CLASP.

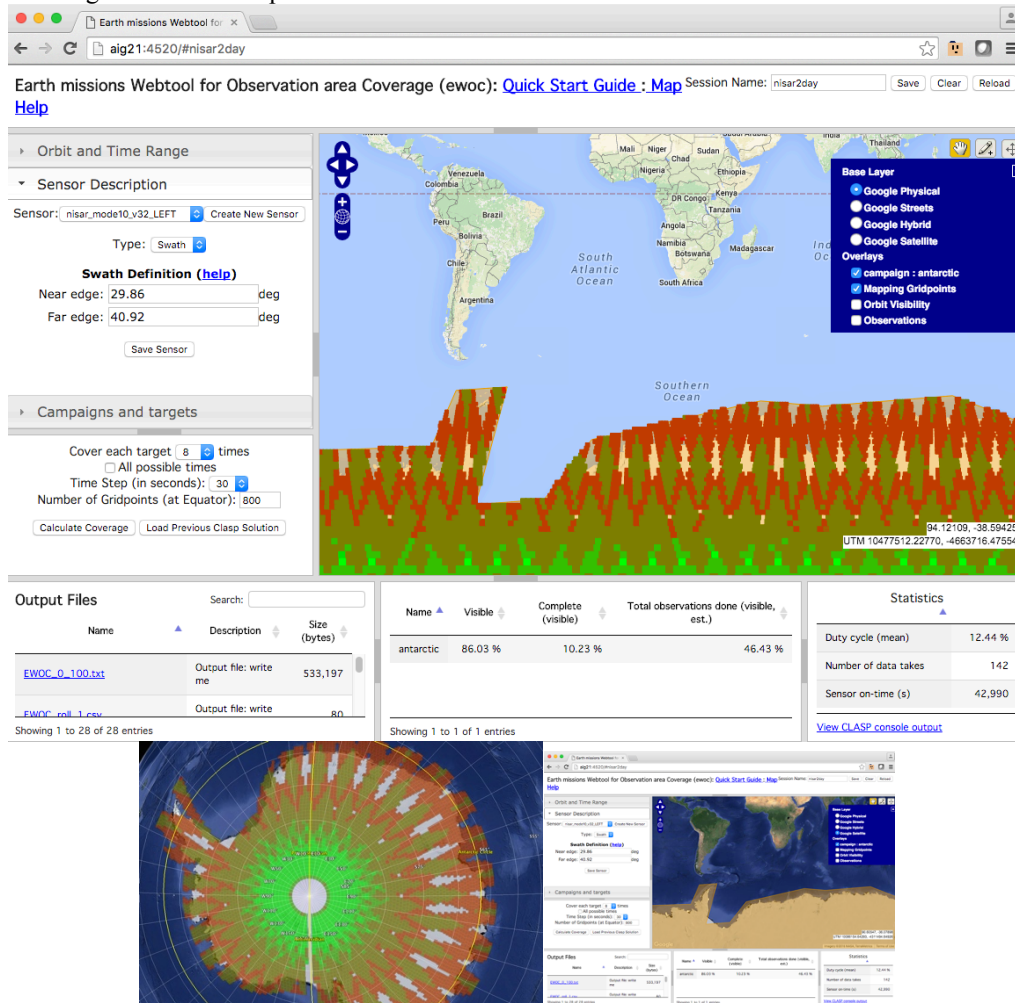


Figure 1. Screen captures of EWOC – a graphical front end for CLASP illustrating the inputs and resulting coverage (top), resulting coverage of 5-day repeat cycle over antarctica in google earth (left) and another view of inputting target polygons in EWOC (right).

In fact, this high-level scope of this study allowed the use of a web-interface to CLASP dubbed EWOC (Earth-missions Web tool for Observation Coverage). EWOC provides a front end for loading pre-generated SPICE kernel files to define the spacecraft orbit, a dialog to define a push-broom style instrument sensor pattern, a geo-spatial tool to quickly define or visualize geographic regions of interest, a widget to define a campaign of observations, and finally an ability to launch the

CLASP tool on a server requesting a single round of scheduling (essentially no optimization) and providing access to download resulting product files and visualization of geographic results. Figure 1 gives an example of a typical view of EWOC with inputs received along the left, and drawn on the interactive map, results along the bottom, and session handling along the top to share and resume work. The same figure is an actual result of a NISAR study; it depicts the expected coverage of a broad region of interest for the cryosphere science discipline over Antarctica using the NISAR instrument in the hypothetical orbit over a five-day period. In particular, this result shows some few small gridded-areas (~10% of total area) receiving all 8 of 8 requested observations over the 5-day period (bright-green) while most of the areas received at least 1 or 2 observations (red to brown), and some few areas receiving no coverage (translucent orange).

These qualitative and quantitative results, in addition to those for a 5-day repeat orbit which while “slower” observed all areas of the target at least once, bolstered positive reviews of the fast-repeat-cycle scenario from the NISAR Science Definition Team’s cryosphere group. High-level plans of navigation to and from these candidate orbits and feasibility of aligning this potential sub-phase of observations into the grander schedule of the mission remains an option to be revisited in later stages of mission development.

## 2.2 Adaptation of Solid State Recorder Model

CLASP has been used extensively to inform decisions toward an SSR architecture and design. Off the shelf CLASP provides a simple monolithic storage model, driven by simple, discrete time-spans of rates of data transfer, which it integrates to a loading profile of the storage device. The resulting profile is exported to a comma-separated text file which can be loaded into a variety of tools from spreadsheet to web-application/libraries (d3.js). This model was used for some time, particularly in pre-phase-A NISAR design and was instrumental in designing the performance merit surface illustrated in [7], as this built-in storage model is a key constraint on CLASP’s optimization loop.

However, designing a space-grade solid state recorder of this capacity (~ 10 Tb) and performance (5Gpbs) pushes today’s technology – a simple monolithic design is not feasible. Proposed designs have varied, but all include some physical partitioning with some ramifications on data IO whether the logical space necessitates partitioning or not. Furthermore, NISAR must maintain certain state information attached to each recorded file of instrument data indicating whether that data is urgent and should be downloaded immediately at next opportunity, and to which agencies’ ground station network the file should-be and has-been downlinked. The NISAR adaptation of CLASP has been developed to include incrementally better fidelity models of proposed SSR designs, leveraging core resource classes from CLASP. To date, the core monolithic storage resource has been extended so as to represent a collection of such devices, each representing either a physical or logical partition and a controller to dictate constraints on how those partitions may be written-to, read-from and otherwise managed as a whole.

While we will not cover specific constraints of various designs, there existed cases with significant complexity in data routing algorithms. For example, routing data to be written to a particular partition could depend upon the state of all the partitions of the system at that particular moment of time in which writing would commence, i.e. writing would occur on the least utilized partition. CLASP however, does not generally schedule operations in a forward sweep of time, it instead, as part of its squeaky wheel optimization algorithm, schedules in dynamic priority order of targets (“squeaky wheel gets the grease” first). Consequently, SSR models with these types of constraints cannot be used in the main scheduling algorithm to limit scheduling decisions and have instead been run as postprocessors. The schedule of observations is instead generated without storage constraints, and merely evaluated against the SSR model after it has been generated.

## 2.3 Observation Strategy

This current strategy of observation planning, while seemingly open loop in a sense, is however leveraging the knowledge gained from the all the previous optimization runs CLASP has run, illustrating many trades between conflicting science campaigns and total data-volume the system is capable of generating and delivering. Together with the Project Scientist and Science Definition Team, the mission planning team has arrived at a very systematic observation strategy in which science disciplines have compromised on mutually acceptable modes of operations to achieve essentially global coverage of solid land-masses and ice. To reduce the amount of data that would be collected from an “always-on-over-land” strategy

in which redundant observations would be made at latitudes approaching the poles where the ground tracks of the

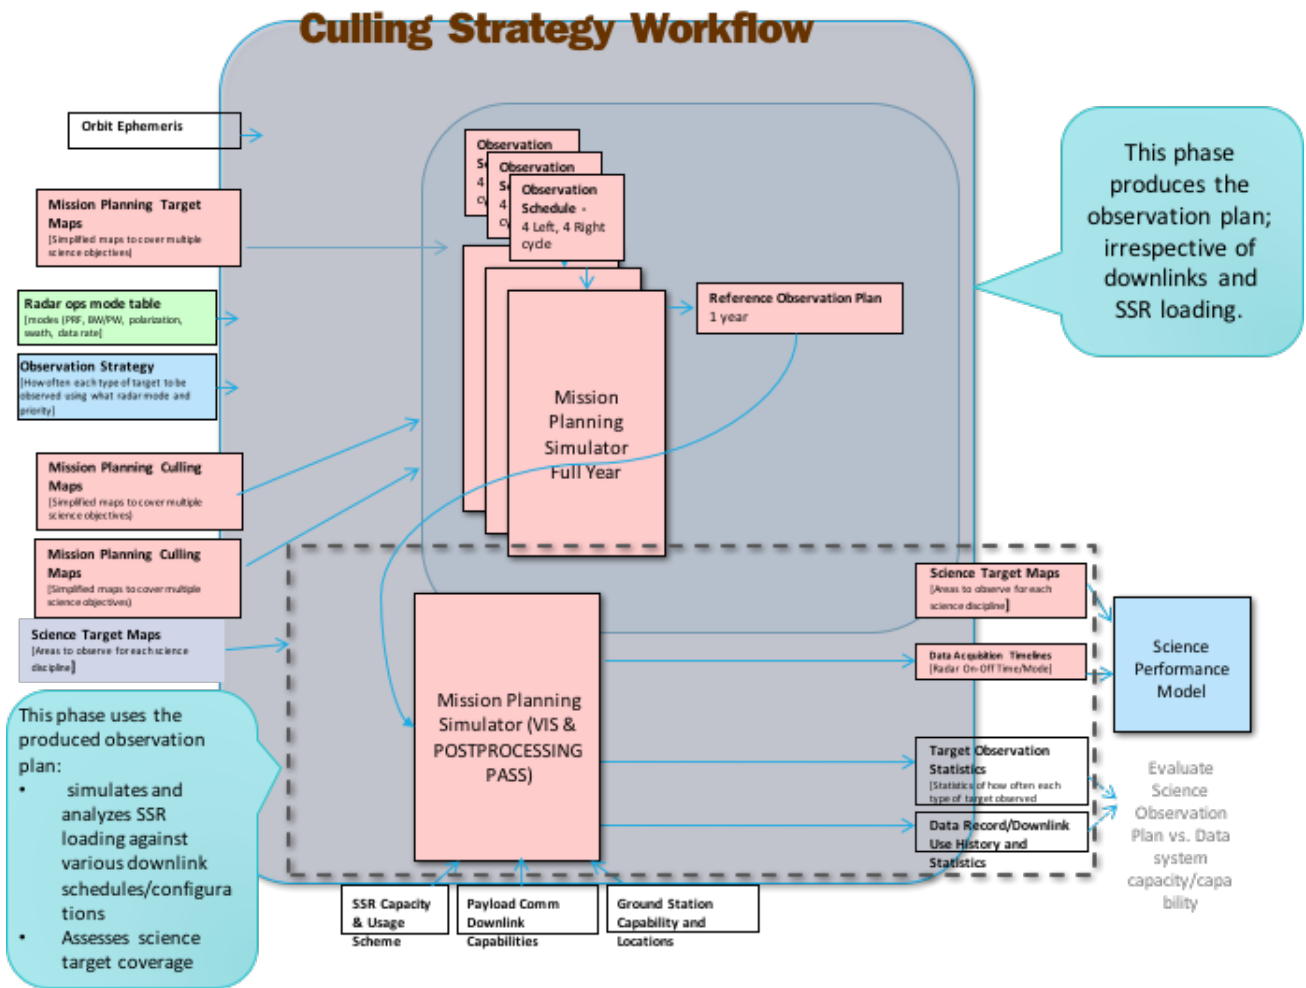


Figure 2. Diagram of flow information in construction of NISAR observation plan.

spacecraft converge, periodic swaths are “culled” from the always-on schedule.

We achieve this culling strategy in CLASP via running it in series with itself and a number of simple filters between. The basic greedy scheduler in CLASP does no backtracking, it cannot remove observations on the current working schedule, it may only add new observations, or promote observations from one mode of operation to a higher, “dominating” mode. We begin the scheduling process by amassing all of the regions of interest, irrespective of their desired mode of operation, and simulate a single sensor mode and swath-width representing the maximum extent of all sensor modes. CLASP digests these inputs with the orbit definition to create a schedule of observations that is essentially all possible viewing opportunities for any target, a baseline schedule. The next invocation of CLASP receives individual science campaigns and an initial priority order, however they are also flagged such that they may not add to the observation timeline, they may only promote over the baseline. This ensures a strict priority ordering – no scheduled science observation may subsequently trump a previously scheduled observation. Each campaign has a desired mode of operation however the same single maximum-extent swath sensor model is used for scheduling observations.

At this stage we have a schedule of observations exceeding (because of the maximum-extent swath sensor model) an “always on over targets” schedule, and we begin our culling strategy. CLASP is run successively with a virtual science campaign “culling” target, a filtering campaign, whose purpose is to promote existing observations to a state in which a post-processor can remove them. We define 3 bands of degree of geographic overlap with respect to each pole of the earth: a triple overlap region, a double overlap region and single-coverage region (including the equator). Within the triple

overlap regions, we keep every third day's worth of observations and cull the remaining two. Similarly, in the double-overlap regions but every other day, while no culling occurs in the single-coverage region. Over a 12-day period these 3-day and 2-day periods will coincide twice, resulting in 2 days in which no culling occurs and long uninterrupted observation tracks are scheduled, a key feature of this strategy. The maps used to define these filtering campaigns for triple and double coverage regions all encompass their respective poles – the double and triple regions overlap so that scheduling observation windows near the geographic boundaries are inclusive rather than exclusive and we default to observing excessively rather than leaving gaps in coverage and missing the real science targets. This last aspect is a recent improvement and has not manifested completely in our baseline reference observation plan. Since these geographic, the 2-day filter could potentially filter out the higher-latitude 3-day region's observations to be kept, however there is an additional filter campaign in priority order between the triple coverage and double coverage filter-out campaigns that essentially flags the remaining triple coverage observations to be kept, i.e. the double coverage filter campaign cannot promote these triple coverage observations to the "to-be-erased" mode.

We have 3 dominant modes of instrument operation at the relevant latitudes that must be culled, a S+L-SAR joint mode for sea-ice observations, and high-resolution L-SAR mode for land-ice, and a L-SAR mode for ecosystems and solid-earth science. A culling target map and filter campaign is created for each mode as they will have subtly different sensor geometries, and each run through a separate invocation of CLASP, with the result from one feeding to the next until we have a culled observation schedule for a 12-day period.

At this point, this 12 day schedule may serve as a template in a process to create an entire year plan, or it may directly be fed into an analysis step. In the analysis step, a table of radar mode parameters is read to create sensor models for every mode used in the observation plan. This is computationally costly as each sensor's swath will be projected over time through the entire planning horizon. Additionally, while the science disciplines have agreed to particular modes of operation and common regions of interest, we still evaluate coverage on more specific target areas – this phase of analysis loads additional target maps, further increasing memory requirements. Recent runs of this analysis step of CLASP have demanded high-water-marks of 20-70GB of RAM, and can go upward depending on runtime options to trade for speed.

### 3. CONCLUSION

#### 3.1 Results

The NISAR mission is currently planning to fly an observation schedule produced through the process described above. Figure 3 below illustrates half of the observation plan projected onto the earth, observations on the descending node of the orbit. On average over a year of operations this plan generates 25.5 Tb of compressed radar data per day, or over 3 petabytes of data over the life of the mission. Duty of the instrument is in excess of 40% and a few observations continue




Figure 3. Illustration of ground swaths of observation schedule on descending node.

back-to-back for over 40 minutes.

### 3.2 Future Work

As the mission continues to solidify in design, so too will the adaptation of CLASP to NISAR. At this time, CLASP does not model “zero-doppler steering”, a common aspect of a radar based spacecraft mission where the attitude of the spacecraft is continuously adjusted to compensate for the rotation of the earth. NISAR will fly such a profile of pointing of the spacecraft, and as such our current model of observations is inaccurate, predominantly timing the observations a few seconds ahead or behind depending on the latitude. CLASP will be modified to input a pointing profile in the form of a SPICE kernel to compensate.

The NISAR SSR model will continue to be developed to capture limitations of file operations. Extremely large files recorded on NISAR can pose a significant overhead to operating the SSR. An improved model will allow trade studies on how to best operate this system, sacrificing observation time for smaller maximum file sizes. The model will also be enhanced to better capture the various states of recorded files and their subsequent playback; whether they have been played-back over an ISRO or NASA station, and whether it is an urgent observation.

#### Acknowledgement

The research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

© 2016 California Institute of Technology. Government sponsorship acknowledged.

### REFERENCES

- [1] Zebker H., “Report of the July 17-19, 2007 Orlando, Florida, Workshop to Assess the National Research Council Decadal Survey Recommendation for the DESDynI Radar/Lidar Space Mission“, NASA Decadal Survey NRC 2007
- [2] Rosen, P. “The NASA-ISRO SAR (NISAR) mission concept,” Pasadena, CA : Jet Propulsion Laboratory, National Aeronautics and Space Administration, 2013, 5-Nov-2013 <[http://trs-new.jpl.nasa.gov/dspace/bitstream/2014/44506/1/13-4891\\_A1b.pdf](http://trs-new.jpl.nasa.gov/dspace/bitstream/2014/44506/1/13-4891_A1b.pdf)>
- [3] Chien, S. et al, “A Generalized Timeline Representation, Services, and Interface for Automating Space Mission Operations,” Proceedings of Space Operations 2012, Stockholm, Sweden, June 2012.
- [4] Choo T., Skura, J., “SciBox: A Software Library for Rapid Development of Science Operation Simulation, Planning, and Command Tools” JOHNS HOPKINS APL TECHNICAL DIGEST, VOLUME 25, NUMBER 2 (2004)
- [5] Choo, T. et al, “The MESSENGER Science Planning and Commanding System” AIAA SPACE 2009 Conference & Exposition 14 - 17 September 2009, Pasadena, California
- [6] Knight R., McLaren, D. , and Hu, S., “PLANNING COVERAGE CAMPAIGNS FOR MISSION DESIGN AND ANALYSIS: CLASP FOR THE PROPOSED DESDYNI MISSION,” International Symposium on Artificial Intelligence, Robotics and Automation in Space i-SAIRAS, 4-6 September 2012, Turin, Italy
- [7] Doubleday, J., Knight, R. “Science Mission Planning for NISAR (formerly DESDynI) with CLASP”, SpaceOps 2014
- [8] Knight, R., Hu, S., “Compressed Large-scale Activity Scheduling and Planning (CLASP)”, IWSS 2009