

# Preliminary Results of Heuristic Guided Orbit Selection for a Low Frequency Radio Interferometric Spacecraft Constellation

Steve Schaffer, Steve Chien\*, Andrew Branch, Sonia Hernandez

Jet Propulsion Laboratory, California Institute of Technology

Pasadena, CA, 91109-8099

firstname.lastname@jpl.nasa.gov

## Abstract

A constellation of radio telescope spacecraft can leverage interferometry to accurately image distant objects throughout the universe, but mission design must balance among many interrelated constraints. In particular, the number of craft and the selection of time-varying orbital parameters play a pivotal role in determining what interferometric baselines are feasible with respect to different targets, and thus drives the breadth and quality of data available to the constellation. The large combinatorial orbit configuration space and competing concerns present a challenging problem that is not well addressed by traditional mission design processes. This paper describes application of automated optimization methods to help direct mission design effort to the most promising dynamic constellation geometries: those that achieve broad interferometric coverage but remain cost-effective and resilient to failures. Several automatic heuristic-driven optimization algorithms representing complementary search strategies were created to explore among concrete constellation configuration plans. Evaluation of each candidate constellation plan was accelerated by efficiently combining precomputed caches of orbital and interferometric data. Results indicate that leveraging automated optimization for constellation mission design is both practical and illuminating: generated solutions provided both evidence for existing design intuitions as well as fresh insights into novel configurations.

## Introduction

Radio-frequency interferometry enables resolving structures of extremely distant objects throughout the universe by combining signal phase data from spatially dispersed sensors (Ryle and Hewish 1960). Imaging capability depends primarily on the variety of target-projected baselines between each pair of detectors (Boone 2002): shorter baselines provide increased field of view, longer baselines yield finer resolution, and relatively rotated baselines reveal oriented spatial structure (Taylor, Carilli, and Perley 1999). A space-based radio interferometer array can usefully extend the advantages of terrestrial arrays in several ways. Foremost, a constellation of radio telescope spacecraft could span enormous baseline distances that are impractical – even impossible – to achieve terrestrially. Additionally, ballistic orbital

motion would continuously vary the array geometry, thus quickly sweeping through a variety of baselines, with even wider possibilities if orbital maneuvers are considered. Finally, a space-borne interferometry array would have almost continuous visibility to targets throughout the sky and in radio frequency bands that are masked by either the Earth's ionosphere or terrestrial noise sources.

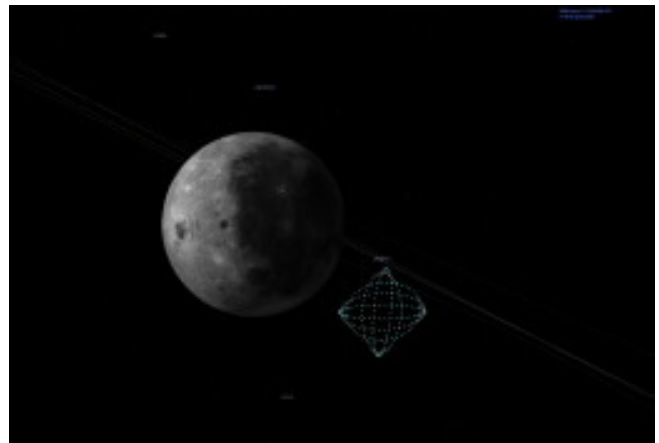


Figure 1: Cosmographia rendering of one possible constellation formation of several concentric elliptical rings in lunar orbit.

Design of the individual radio telescope spacecraft and the array as a whole must be carefully balanced to leverage these advantages against the increased material and logistical costs of deployment in space. Typical spacecraft constraints such as launch mass, delta-v budget, communication capability, operational complexity, failure resilience, and mission cost are magnified simply by the multiplicity of craft in a constellation, but interferometry introduces its own additional obstacles. Each radiometer element could produce data rates up to several megabits per second, and this raw data must be correlated with high precision timing and position data in order to allow accurate interferometric reconstruction. Scientifically valuable long baselines between the detectors increase the equipment and power costs to relay data among the craft, with eventual transmission to Earth likely handled by a subset of craft with sufficiently powerful transmitters. Individual constellation craft failures are also

\*corresponding author

relatively more costly, since they degrade the baseline density of the entire array quadratically rather than linearly. And finally, the number of craft and the individual orbits must be jointly selected from a huge combinatorial variety of possibilities in order to maximize the science data quality for targets widely distributed throughout the sky.

Launching an array of radio astronomy telescopes has been studied in the past (Rajan et al. 2016), though no such mission has flown to date. Contrary to the detailed analysis of ground-based interferometer layout (Keto 1997), existing constellation mission designs typically sketch out only very high level (eg spherical) array geometry without detailing how individual orbits for each craft should be selected to maximize mission science value. Since array geometry is a key driver for many constellation performance measures (eg: image quality, overall cost, fuel budget, and failure resilience), the orbital selection process should be the target of aggressive engineering. This work serves to fill in the gap between design intuition and quantitative optimization by providing an automated analysis approach to study the trade-offs of different orbit geometry choices. A model mission scenario is described and then used to drive a series of efficient heuristic-guided optimization algorithms in order to both confirm existing design intuitions and synthesize new ones. Furthermore, the heuristic-based optimization approach is highly portable and could be leveraged in a variety of mission contexts to help focus human design effort.

### Mission Concept

The present mission study (Wyatt et al. 2017) posits a constellation in lunar orbit with star topology: a single communications relay mothership flanked by a few dozen radio telescope microsatellites. A single launch vehicle first delivers the mothership assembly into a lunar orbit, which then deploys the detector craft into concentric rings with varying inclinations. The daughter ships carry additional fuel for station-keeping and later orbit modifications in which an entire ring grows to different discrete size stages. Each detector craft uses a twin-dipole antenna to acquire a full-sky analog signal that is then digitally sampled, time tagged, and compressed locally before eventually being transferred back to the mothership for relay to Earth. Computing resources on the ground uncompress and interferometrically recombine the phase data from each craft to generate imagery of various radio sources, exemplified by active galactic nuclei.

This operational scenario admits a huge variety of possible constellation configurations, with the foremost design variable being the number of detector craft in the constellation, followed by the ring organization structure and geometric parameters (Hernandez et al. 2017). Our results explore among concrete scenarios chosen from a cartesian product of representative discretized ring parameters:

- $N \in \{2 \dots 70\}$ , the total number of detector craft in the constellation, across all rings
- $R \in \{1 \dots 28\}$ , the number of rings that the detector craft are organized into
- $n_r \in \{1 \dots 10\}$ , the number of craft in a specific ring  $r$

- $a_r \in \{100, 300, 500 \text{ km}\}$ , the semi-major axis of the ring around the mothership, which may increase via discrete boosting maneuvers during the mission
- $i_r \in \{0, 15, 30, 45, 60, 75, 89^\circ\}$ , the inclination of the ring above the orbital plane of the mothership
- $\nu_r \in \{0, 36, 72, \dots, 324^\circ\}$ , the angular position of the craft within the ring at the epoch time, similar to orbital anomaly

The range of actually valid configurations is limited by the interdependence of constellation-wide fuel mass available, boosted ring sizes, and ring inclinations. Even with reasonable discretization, the combinatorics of the solution space are exponential in the number of craft, and quickly become prohibitive for exhaustive examination even for relatively small constellations. Instead, our approach locates high quality array orbit plans (that may include ring size boosts) using several heuristic-guided search methods.

### Approach

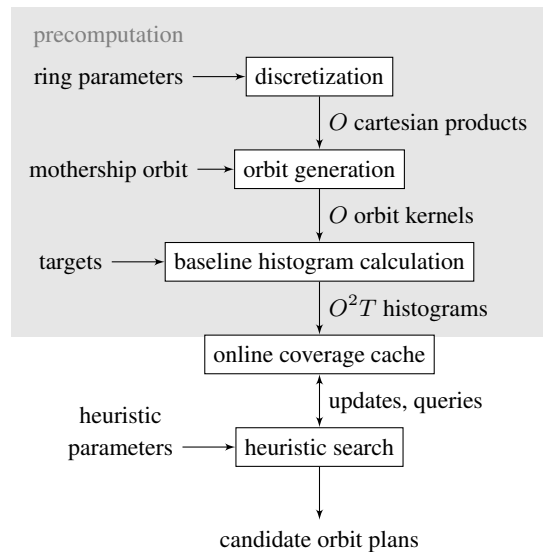


Figure 2: Process data flow outline, showing offline preprocessing steps that increase efficiency of subsequent optimization runs.

The objective of the automatic orbit optimization process is to present the human mission designers with several viable and high quality constellation orbit plans for further study. Each output constellation orbit plan consists of a concrete choice of ring parameters for each spacecraft throughout the mission, including possible ring size boosts. The results need not be strictly optimal, but should represent a variety of promising candidates from the vast solution space. Solution quality is judged first according to adherence to specified mission constraints (such as limited total mass or number of craft), and then via a weighted sum of objective functions (such as maximized image quality, maximized failure resilience, or minimized total propellant use). Several different pluggable heuristic search strategies are implemented

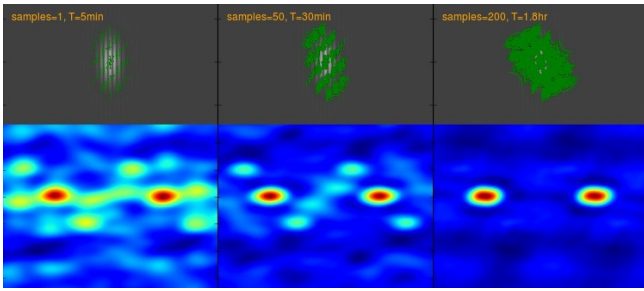


Figure 3: Successive output from APSYNSIM interferometry simulation, showing (top) increasing interferometric baseline coverage for a given target over the course of a mission, and (bottom) corresponding improvement in reconstructed image quality.

within this general framework, which also provides for convenient comparison and composition of the individual algorithms. To achieve reasonable efficiency, our approach strategically discretizes the parameter space and pre-caches complex orbital and interferometric calculations, while also aggressively parallelizing each evaluation step. The key enabling insight is that the paramount image quality metric is primarily driven by the variety of target baselines collected. Thus, the search algorithms can keep a running tally of baseline coverage in an efficient histogram structure rather than running a fresh interferometric simulation at each step in the state space search.

The processing pipeline begins with an offline precomputation of relevant orbits and interferometric baselines in order to completely fill a cached set of baseline histograms used during active search. The first step is a user-guided discretization of the ring parameters, which yields a cartesian product of candidate parameterizations sampling the continuous configuration space. Then, the orbit parameterizations are combined with the mothership’s lunar orbit to generate a set of differential equations for each daughter spacecraft’s position and velocity as a function of time (Hernandez et al. 2017). A matlab script numerically solves these equations of motion and then samples them over the course of a full lunar orbit cycle, yielding time series of state vectors. In turn, these state vectors are transformed into efficient general purpose NAIF SPICE compatible `.bsp` orbit kernels by the the SPICE `mkspk` tool (Acton 1996).

Next, all of the baselines between potential members of the constellation are computed with respect to each of a set of specified celestial targets. In practice, the intended targets are expected to be isotropically distributed throughout the sky, and so a synthetic set of 22 targets that evenly sample the celestial sphere is used in this study. The SPICE kernels from above are used to compute a time series of projective interferometric baselines by sampling spacecraft and target geometries over the course of a full orbit cycle. The baseline measurements are condensed into a single summary histogram for each of the possible orbit-orbit-target tuples, using a specified histogram binning. A reasonably-sized polar histogram with 100 radial bins by 100 angular bins, uniformly spanning  $r \in [0, 700 \text{ km}]$  and  $\theta \in [0, \pi]$  was used in this study. The histograms are then written in compressed

form to a large cache file on disk, from which they can be readily reconstituted on demand during active search.

Each histogram cache line represents the interferometric coverage of a given target provided by a pair of hypothetical spacecraft orbits that could eventually be part of the constellation. Thus, the individual histograms can be additively combined for all concurrent pairs of orbits in any proposed constellation plan to estimate the baseline coverage of the whole constellation. This fast incremental update of evaluation scores is particularly advantageous for local search strategies that require efficient heuristic guidance, and is also very amenable to parallelization.

Several heuristics were implemented that leverage the precomputed interferometric baseline coverage histograms. The most important is an image quality estimator based on the portion of non-empty (ie covered) histogram bins. A loss-sensitive heuristic measures the expected value of the same coverage metric proxy in the face of some specified risk of craft loss. Further heuristics were implemented in order to score solutions based on their propellant mass requirements: both to attain and maintain desired ring inclination, as well as to boost ring sizes throughout the mission plan. The heuristics can be combined with each other under a weighted sum in order to balance solutions among competing concerns simultaneously.

The above heuristics are used by a variety of optimization algorithms to productively explore the possible constellation configuration plans. The first algorithm simply randomly selects a given number of candidate spacecraft to include in the constellation, including growth stages for each ring from the minimum to maximum size. Since the random algorithm does not conduct any guided search, several random trials are averaged, and the results are presented mainly as a point of comparison.

The first local search option is the forward greedy algorithm. Starting from the empty constellation plan, it incrementally adds one new orbit stage to the plan (either by adding a new craft entirely, or appending an additional growth stage to an existing ring). At each iteration, the search algorithm evaluates (in parallel) all permissible options for the next step, and the option with the highest heuristic value is selected as the basis for the next iteration. The process repeats until a termination criteria is reached (typically a specific target number of craft in the constellation). A symmetric reverse greedy algorithm was also implemented. It starts with an exhaustively populated constellation plan including all possible craft and ring growth stages, and then incrementally removes a single orbit stage from the plan that minimizes the detriment to its heuristic score. A natural extension of both greedy algorithms begins with a randomly constructed plan and then explores in the vicinity of that starting point.

The two single-direction greedy search algorithms were also combined into a bidirectional greedy search algorithm. Starting from some (possibly randomized) initial state, the bidirectional algorithm either adds or removes an orbit stage from the constellation plan at each step. Each iteration has a choice of both which direction it will move the solution (larger or smaller), and also how it will choose the element to

insert/remove (randomly, greedily, or otherwise). The iteration choices are provided as a vector input to the algorithm, allowing for a kind of simulated annealing by alternately relaxing and refining a working solution.

### Initial Results

The randomized constellation algorithm achieved surprisingly good baseline coverage scores for a variety of constellation sizes, although it disproportionately suffered at low sizes. This is explained by the natural combinatoric tendency of sufficiently large random solutions to sample evenly from a variety of ring planes, thus providing for a good chance at a variety of baselines perspectives. However, at lower constellation sizes of interest the actual choice of which ring planes to populate becomes critically important, and the blind random algorithm correspondingly underperforms.

Both deterministic single-direction greedy algorithms out-perform the random algorithm mean in baseline coverage along the entire spectrum of constellation sizes by exploiting domain knowledge encoded in the heuristic function. The benefit of the guided search approach is especially magnified at small constellation sizes. The reverse greedy algorithm achieves matched size constellation solutions with better coverage than the forward greedy algorithm; a welcome trade for the more extensive computation time expended to descend meticulously from the initial exhaustive constellation. Indeed, the reverse greedy algorithm starts from a position of superior contextual information and must simply select a modification that minimally damages its already excellent candidate solution. Conversely, the forward greedy algorithm must start from nothing, rendering it much more sensitive to any initial missteps along the search trajectory.

to mitigate this kind of sensitivity, but actually degrades the performance of the forward greedy algorithm for small constellations (since a relatively large portion of the solution is now purely random). The random initial starts only marginally improve the final coverage scores from the reverse greedy algorithm, but they do vastly reduce the execution time of the algorithm by skipping long chains of nearly-tied craft removal decisions. Adding simulated annealing behavior to the algorithms should also reduce the lock-in effect of early decision points by occasionally relaxing solutions away from the final target size. Such a bidirectional search strategy appears to eke out a small incremental gain in final coverage score for the reverse greedy algorithm under finely tuned annealing magnitude and frequency parameters.

The above algorithm solutions are closely optimized and are thus susceptible to loss of key members of the constellation. Introducing the loss-sensitive heuristic is very successful at mitigating this kind of loss sensitivity, and leads to solutions that exhibit graceful degradation in the face of predicted craft losses. The added robustness comes in the form of built-in redundancy in the constellation, and thus hurts the coverage performance of the constellation in scenarios where no losses actually occur, but pays dividends when losses do occur. Interestingly, the benefit is quite sensitive to the predicted loss rate provided to the heuristic: at actual loss rates well above or below the predicted rate, a mismatched robust solution provides no benefit over the purely greedy solutions.

### Conclusion

Heuristic guided optimization is usefully applied to the problem of orbit selection for a constellation of radio interferometric spacecraft. The primary driving feature of such a constellations is its image reconstruction capabilities, which while not presently efficient to directly compute, has a satisfactory proxy in baseline histogram coverage. Carefully orchestrated precomputation and parallelization steps allow the baseline histograms for intermediate solutions to be computed very efficiently, vastly speeding the execution times for optimization algorithms relying on coverage heuristics.

Automated optimization processes guided by simple heuristics allow mission designers to explore and evaluate vast mission configuration spaces with relative ease. Automation of otherwise tedious design choices allows designers to focus their efforts on higher quality solutions. In addition, the generated solutions and analytic metrics provoke fresh insights into the problem that are portable to future design iterations. These improvements in the design process could lead to increased productivity and cost savings, along with improved final mission science value.

### Acknowledgments

Portions of this work were performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

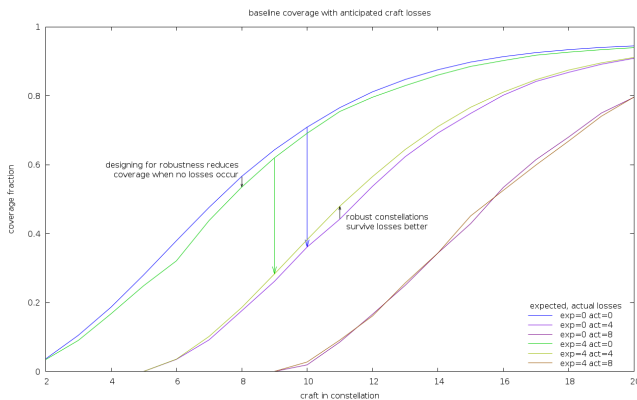


Figure 4: Explanatory plot of coverage metrics (out of 1.0) for solutions generated by using loss-sensitive optimization heuristic compared with the default heuristic. In situations with no realized losses, the loss-sensitive solution (green) fares worse than the naive solution (blue). However, when the predicted loss rate actually occurs, the loss-sensitive solution (yellow) dominates the naive solution (purple). At even higher loss rates, the solution difference becomes equivocal.

Adding random initial starts to each algorithm attempts

## References

- Acton, C. H. 1996. Ancillary data services of nasa's navigation and ancillary information facility. *Planetary and Space Science* 44(1):65–70.
- Boone, F. 2002. Interferometric array design: Distributions of fourier samples for imaging. *Astronomy and Astrophysics* 386(3):1160–71.
- Hernandez, S.; Garza, D. M.; Broschart, S. B.; Herzig, S. J. I.; and Chien, S. A. 2017. Small satellite constellation to enable a lunar radio interferometer.
- Keto, E. 1997. The shapes of cross-correlation interferometers. *The Astrophysical Journal* 475(2):843.
- Rajan, R. T.; Boonstra, A.-J.; Bentum, M.; Klein-Wolt, M.; Belien, F.; Arts, M.; Saks, N.; and van der Veen, A.-J. 2016. Space-based aperture array for ultra-long wavelength radio astronomy. *Experimental Astronomy* 41(1-2):271–306.
- Ryle, M., and Hewish, A. 1960. The synthesis of large radio telescopes. *Monthly Notices of the Royal Astronomical Society* 120:220.
- Taylor, G. B.; Carilli, C. L.; and Perley, R. A., eds. 1999. *Synthesis Imaging in Radio Astronomy II*, ASP Conference Series.
- Wyatt, E. J.; Castillo-Rogez, J. C.; Chien, S. A.; Clare, L. P.; Fraeman, A. A.; Herzig, S. J.; Nesnas, I. A.; and Lazio, J. 2017. Novel planetary science enabled by networked constellations. In *Planetary Science Visions 2050 Workshop*.