The Astronomical Journal, 172:288 (13pp), 2026 December https://doi.org/10.3847/1538-3881/aeaa59
© 2026. The Author(s). Published by the American Astronomical Society.
OPEN ACCESS

A Large-scale Technosignature Survey of the Antisolar Point with the Allen Telescope Array: Survey Design and Pilot Results

Sofia Z. Sheikh[1,2], Eleanor Hort[3], Bárbara Cabrales[4], Isabel Gerrard[1,6], Wael Farah[1], Nick Tusay[5], David R. DeBoer[6,7], Saul Balcarcel-Salazar[8], Joe S. Bright[1], Luigi F. Cruz[1], Ben Jacobson-Bell[2,9], Mark Masters[1], Alexander W. Pollak[1], Seth Shostak[1], Carol Shumaker[1], and Andrew Siemion[1,2,4,6,10]
[1] SETI Institute, 339 Bernardo Ave., Suite 200, Mountain View, CA 94043, USA
[2] Berkeley SETI Research Center, 339 Campbell, Berkeley, CA 94720, USA
[3] Department of Physics and Astronomy, University of New Mexico, 210 Yale Blvd. NE, Albuquerque, NM 87106, USA
[4] Jodrell Bank Centre for Astrophysics (JBCA), Department of Physics & Astronomy, Alan Turing Building, The University of Manchester, M13 9PL, UK
[5] Department of Astronomy, University of Washington, Seattle, WA 98195, USA
[6] Sub-Department of Astrophysics, University of Oxford, Banbury Road, OX1 3RH, UK
[7] Radio Astronomy Laboratory, University of California, Berkeley, CA 94720, USA
[8] Department of Physics, Stanford University, 382 Via Pueblo Mall, Stanford, CA 94305, USA
[9] Department of Astronomy, University of California, Berkeley, 501 Campbell #3411, Berkeley, CA 94720, USA
[10] Department of Physics, Faculty of Science, University of Malta, Msida MSD 2080, Malta
Received 2025 November 6; revised 2026 September 3; accepted 2026 September 9; published 2026 November 3

Abstract

In the search for technosignatures, it can be useful to narrow down the cosmic haystack by prioritizing time-coordinate "Schelling points" at which a beacon may be more likely to be transmitted. One such Schelling point is the antisolar point (ASP), the place and moment at which Earth would lie directly between a hypothetical transmitter in the Earth's transit zone (i.e., ecliptic) and the Sun. In this work, we describe a campaign to survey the ASP for radio technosignatures using the Allen Telescope Array, report the results of 3 initial nights of data from a pilot study in summer 2023, and discuss the future of the campaign. In the pilot study, we find no sky-localized narrowband 1–2.4 GHz technosignatures within an R.A. range of approximately 20.266–21.312 hr during the ASP, setting an upper limit flux density threshold of 115–936 Jy.

Unified Astronomy Thesaurus concepts: Search for extraterrestrial intelligence (2127); Technosignatures (2128); Ecliptic (2336); Radio interferometry (1346); Radio spectroscopy (1359); Astrobiology (74)

1. Introduction

1.1. The Antisolar Point

"Technosignatures," astronomically observable phenomena generated by technologically capable life, provide a complementary astrobiological search strategy to "biosignatures" in the search for life in the Universe (J. T. Wright 2020). The microwave frequencies have been historically favored as a place to search for technosignatures, with a strong scientific basis well-described in the frame of the nine axes of merit for technosignature searches (S. Z. Sheikh 2020). Radio waves traverse the galaxy at the speed of light, are more energy-efficient in bit rate (especially if they have a narrow bandwidth) than higher-energy photons, and are not attenuated by galactic gas and dust. In addition, the astronomical radio sky is generally quite quiet, with only a few morphologically distinct classes of natural confounders that look very different than the signals produced by human technology; human radio transmitters tend to emit signals that are orders of magnitude spectrally narrower than masers, the narrowest natural phenomena (J. Tarter 2001). Finally, using radio technology to communicate is something that humans have done for approximately a century, meaning that there is no extrapolation needed to unknown technologies.

In technosignature searches, we often make an assumption that we are looking for an intentional communication that is being broadcast by an extraterrestrial agent (ETA) to interstellar space.[11] For radio searches specifically, this assumption is broadly driven by the fact that even the strongest omnidirectional "leakage" radiation from the Earth (i.e., 4G LTE transmissions; R. C. Saide et al. 2023) is thousands of times weaker than our strongest directional transmission capabilities on Earth (R. D. Ekers et al. 2002).

Even after the simplified assumption of an intentional beacon, there is still an immense eight-dimensional parameter space to search (J. T. Wright et al. 2018). Thus, we additionally assume that energy conservation and efficiency are important to the ETA sending the signal. There are two general ways to conserve energy in interstellar radio signaling, assuming a fixed morphology: (1) send signals only intermittently, instead of persistently, and (2) beam the power in a specific direction (instead of building an omnidirectional transmitter) (R. H. Corbet 2003). The issue with these low-duty-cycle, beamed signals is that they are unlikely to be picked up from sky surveys by chance; while upcoming all-sky, all-the-time programs such as PanoSETI (J. Maire et al. 2022) may have a chance of serendipitously observing them, no such programs exist today, and none exist at radio wavelengths.

To mitigate the detectability issue with beamed, low-duty-cycle transmissions, we can rely on the idea of mutually

Original content from this work may be used under the terms of the Creative Commons Attribution 4.0 licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

[11] Some recent papers avoid this assumption, such as the planet–planet occultation search in N. Tusay et al. (2024).

determinable "Schelling points" (T. C. Schelling 1980) from the game theory literature, which has made interdisciplinary inroads in the technosignature literature (J. T. Wright 2020). To apply this theory to an ETA's beamed transmission to Earth, we need to consider which astronomical parameters can be used to signal a mutually determinable *time* and *coordinate*. Some authors have proposed Schelling points in time synchronized to supernovae (P. Makovetskii 1980; J. R. Davenport et al. 2022; B. Cabrales et al. 2024), nearby variable stars (W. I. McLaughlin 1977), or gamma-ray bursts (R. H. Corbet 1999).

Here, instead, we wish to consider a characteristic of our own solar system visible at astronomical distances, with the assumption that Earth is the intended recipient of the signal. If the transmitter happens to be located in the Earth's transit zone (ETZ), defined by a narrow range of angles above and below the ecliptic, then the ETA would be able to see the transit of the Earth across the face of the Sun (T. Castellano et al. 2004). Transits are an excellent feature to use for a Schelling coordinate because they are easily detectable, they provide extra information about our planetary system (e.g., the existence of a planet in the habitable zone, opportunities for planetary atmospheric characterization), and they give clear opportunities for temporal synchronization (S. Z. Sheikh et al. 2020).

With the Earth's transit as our chosen feature, one potential Schelling time is defined when the Sun is directly opposite the transmitter in the sky, i.e., when the transmitter sees Earth at the midpoint of its transit: from here on, we shall refer to the apparent position (from Earth) of this alignment as the antisolar point (ASP). This idea was first articulated, for Jupiter's transits specifically, by C. E. Singer (1982), but Earth may be preferable due to its transit-derivable presence in the habitable zone (R. H. Corbet 2003). The ASP Schelling coordinate/time is generated by a simple geometry with no obvious degeneracies, and it is related to the target habitable planet (Earth) directly, with no need to rely on external (as a nearby periodic astronomical system) or unpredictable (as a high-energy explosion) events (R. H. Corbet 2003).

However, there are some caveats to the ASP method. First, either the transmitter or the receiver must account for the light travel time that will cause Earth's midtransit to appear at different times to an observer on Earth and a distant transmitter. As the transmitter, the ETA would have to know the exact distance and proper motion of the Sun, including the orbit and orbital phase of the Earth (R. H. Corbet 2003). S. Shostak & R. Villard (2004) argue that the number of targets in the transmitter's survey is very likely much smaller than the number of potential near-ecliptic stars as seen from Earth, so the onus should be on the transmitter to correct. Given that the ETA that is transmitting will probabilistically have arisen much earlier than humans did (i.e., D. Kipping et al. 2020), they may have better technology for planet detection than we do on Earth.[12] The critical factor, the distance to the system, could be determined to high precision with a more sophisticated version of a parallax mission analogous to Gaia (T. Prusti et al. 2016). Note that this survey is also sensitive to probes within the solar system (R. Freitas 1980), from which it is easy to calculate the required time delay.

[12] Of course, this improved technology could also reduce their reliance on the transit method entirely.

Despite this strong motivation for an ASP search, no radio technosignature campaign has surveyed with the necessary coordinates and timing. As of 2020, the Breakthrough Listen Initiative (S. P. Worden et al. 2017) did not have a single pointing within 0.°5 of the ASP within their Open Data Archive.[13] Large radio technosignature surveys such as META (F. Colomb et al. 1995) and BETA (D. Leigh & P. Horowitz 2000) likely spent only 0.0004% of their survey time on the ASP, entirely serendipitously; the optical search performed by P. Horowitz et al. (2001), at best, would have resulted in 0.05% (R. H. Corbet 2003). A dedicated survey will be necessary to place limits on radio technosignatures leveraging the ASP as a Schelling point.

1.2. The Allen Telescope Array

The Allen Telescope Array (ATA) is a 42-dish radio interferometer located at Hat Creek Radio Observatory (HCRO) in Hat Creek, California. It is operated by the SETI Institute and was purpose-built for radio technosignature surveys (J. Welch et al. 2009), making it the optimal instrument for a long-term ASP survey.

The ATA, as a ground-based facility, can only observe the ASP when it is above the horizon at Hat Creek, i.e., during the nighttime, which would lead to a total survey coverage of approximately half of the ecliptic over 1 yr (more densely populated during the Northern Hemisphere winter months and more sparsely populated during the Northern Hemisphere summer months). Due to the ∼365.25 day orbit of the Earth and the presence of a leap day every 4 yr, the available swath of the ecliptic as viewed from Hat Creek precesses, such that over 4 yr, the ASP will be observable over the entire ecliptic. Covering this entire region would result in a sky coverage improvement of 4 orders of magnitude for this technique.

Each 6.1 m element of the ATA is being retrofitted with cryocooled, log-pyramidal "Antonio" feeds, which are sensitive from 1 to 10 GHz. Currently, 28 of the 42 antennas have been outfitted with the new feeds and are used as part of the high spectral resolution beamformer backend that is being deployed on the ASP survey. Within the 1–10 GHz range, the science user can tune the on-site local oscillators to any four 672 MHz spectral windows for observation; note that only two were operable until summer 2025, and the initial survey design maintains the two-tuning setup.

In Section 2, we will discuss the observing strategy, analysis strategy, and upper limits that will be set by the ATA-ASP survey. In Section 3, we discuss the results from the 3 nights of pilot observations with the ATA in summer 2023, setting some initial limits on a few discrete ASP pointings. Finally, we will discuss the future of the ATA-ASP survey and conclude in Section 4.

2. Survey Design

2.1. Observing Strategy

If the ATA were dedicated to ATA-ASP for 100% of its observing time, the available half of the ecliptic would be covered exactly 1 yr after it took its first night of data. Two years later, the campaign could be reperformed to cover the

[13] https://breakthroughinitiatives.org/opendatasearch; this work was performed by coauthor Balcarcel as a final project for the online Introduction to Astronomy Research course, advised by coauthor Sheikh: https://github.com/howardisaacson/Intro-to-Astro-2020.

other half of the ecliptic. However, other projects will also be in the nightly observing queue, especially priority sources for other ATA science areas such as fast transients (e.g., fast radio bursts, FRBs; S. Z. Sheikh et al. 2024) and slow transients (e.g., supernovae; J. S. Bright et al. 2023). Given that, we expect ATA-ASP to be on-sky approximately 2–3 days week^{-1} and for the survey to therefore take about 4 yr to obtain full coverage of the ecliptic.

For this survey, we are utilizing the ATA's high spectral resolution beamformer, which runs on the `Breakthrough Listen Accelerated DSP Engine (BLADE)` software.[14] The beamformer produces a `filterbank` HDF5 "data cube" of time, frequency, polarization, and flux data corresponding to particular points within a single dish's field of view. Note that we do not perform any coherent Doppler correction (e.g., to the solar system barycenter) before production of the data cube but instead incoherently correct for drift rates in the `bliss` stage of the analysis pipeline (Section 2.2.1). Coherent de-Doppler methods allow for more power to be recovered for drifting signals, as discussed in J.-L. Margot et al. (2021), but are computationally expensive; we plan to develop and deploy de-Doppler methods on ATA eventually, but they are not used in this work.

Technosignature searches can, in theory, be conducted in the image plane as well, giving certainty to the localization of emission on the sky, but that strategy is still under development for ATA data. To ensure sky localization in a beamforming mode, we instead form two simultaneous synthesized beams on the sky: one at the location of the ASP and one a suitable number of synthesized beamwidths away but still within the full width at half-maximum (FWHM) of the primary beam (as employed by N. Tusay et al. 2024). This is similar to the "nodding" or ABABAB strategy used by single-dish technosignature searches such as C. Choza et al. (2023). For an array, signals that (a) appear in the target "on-beam" and (b) are significantly attenuated in the control "off-beam" are candidates for signals that are sky-localized: distant enough from Earth that they appear as point sources.

We chose an observation time of 5 minutes for each coordinate pointing on the ASP. This duration is short enough that the ASP's precession should not significantly affect its position within the beamformed beam but long enough to be able to distinguish drifting signals. It also creates output files that are manageable in data volume, despite requiring extremely fine sampling in frequency. The frequency resolution in these data is 1.9 Hz, with a time resolution of 16.8 s. The 1.9 Hz sampling is sufficient to resolve narrowband (Hz-scale) signals; combined with the 16.8 s integration time, physically meaningful drift rates are easily identifiable in the dynamic spectra.

As described in Section 1.2, the ATA is capable of recording up to ∼2.8 GHz of simultaneous bandwidth, divided into four independently movable 672 MHz chunks. For this survey, to keep data rates manageable with the high observing cadence, we select and record with only two of these tunings. While the *L* band (1–2 GHz) was favored in early Search for Extraterrestrial Intelligence (SETI) experiments (because it contains the "Water Hole"; R. E. Machol et al. 1975), modern radio technosignature searches have been expanding through the radio spectrum, with increasing numbers of observations taken in the 1–12 GHz range by the Breakthrough Listen Initiative (BL) (e.g., S. Z. Sheikh et al. 2020). If we assume a uniform prior across frequencies (not relying on "magic frequencies"; B. Zuckerman 1985) and note that no frequency has been thoroughly explored as of yet in a parameter space or "cosmic haystack" sense, then practical considerations will dictate our choice of frequencies. The two practical considerations that drove our frequency selection were (1) beam size and (2) radio frequency interference (RFI).

R. Wells et al. (2018) estimate the ETZ to have a width of 0.°53. The ATA has a large primary beam size, such that the primary beam covers the entire ASP for any frequency below 6.5 GHz. However, the synthesized or beamformed beamwidth is smaller than the primary beamwidth by a factor of approximately 50—there is no ATA-observable frequency for which the synthesized beam will cover the entire ASP. At such high frequency resolutions and, commensurately, data rates, the ATA can currently record up to two simultaneous synthesized beams. Thus, to cover as much of the ASP as possible, lower frequencies (resulting in larger synthesized beams) are preferred.

However, the second consideration is RFI, which is notoriously bad (especially at the *L* band) and only worsening (e.g., L. Schirru et al. 2023). To work around the worst regions of RFI contamination at HCRO while still keeping the synthesized beam as large as possible, we selected center frequencies of 2692 and 3364 MHz, resulting in a frequency range of 2356–3700 MHz.

To further improve the robustness of the survey to RFI from ubiquitous satellite communication constellations, we have implemented telescope boresight avoidance (TBA; B. D. Nhan et al. 2024). TBA is implemented by sharing the observing location with the satellite operator (in the first instance, the direct-to-cell service enabled by the SpaceX Starlink satellites), and they moderate the transmission by either limiting beamforming or turning off the transmitter as the transmitting satellite is near or traversing the main beam, respectively. The direct-to-cell signal occupies the band of 1990–1995 MHz, which is outside of the band in this work; however, the signal strength is such that without this technique, the entire RF band of the ATA is saturated and unusable during the roughly 15 s of the traverse. Currently, for a 10 hr observation, TBA is applied to mitigate direct-to-cell in, on average, 50 independent instances.

The geometry of the observing scheme is shown in Figure 1. We place one synthesized beam on the center of the ASP (aka the region directly on the ecliptic), which corresponds to observers that see the Earth transit the Sun with the lowest impact parameters. This covers the innermost ∼5% of the ASP's diameter: 0.°030 at 2356 MHz and 0.°019 at 3700 MHz. The other synthesized beam is reserved for the off-beam and is located five synthesized beamwidths (at the lowest frequency in the band) from the on-beam, perpendicular to the ecliptic. Five beamwidths at the lowest frequency in the survey is equivalent to ∼eight beamwidths at the highest frequency in the survey; this is large enough to get isolation between the two synthesized beams at the lowest frequencies but also small enough that the off-beam will still be comfortably located inside the primary beam and the ASP even at the highest frequency in the band. As the ASP sweeps across the ecliptic, target stars enter and leave the region (horizontally, in the

[14] The software is available on the GitHub `BLADE` codebase, https://github.com/luigifcruz/blade, under an MIT license. The version used for this analysis is archived in Zenodo (L. Cruz 2024).

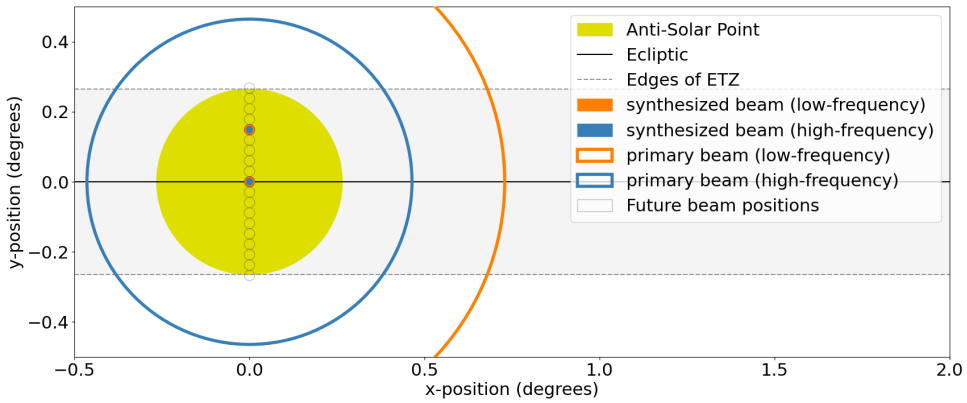


Figure 1. The on-sky geometry of the ATA-ASP campaign. The black horizontal line represents the ecliptic, with the gray shaded region indicating the edges of the ETZ. The size of the ASP at any one moment is displayed in yellow, with the lowest- and highest-frequency primary beam sizes overplotted for scale (red and blue circular outlines, respectively). The synthesized beam sizes and positions for the lowest and highest frequencies are shown as shaded red and blue regions, respectively, within the ASP. As the survey continues, we may add more beams in the positions indicated by the gray circular outlines within the ASP or shift the survey to the image plane, where the imaged region would be represented by the primary beam sizes. It should be noted, however, that the initial two-beam survey is still covering a fully uncharted region of technosignature parameter space.

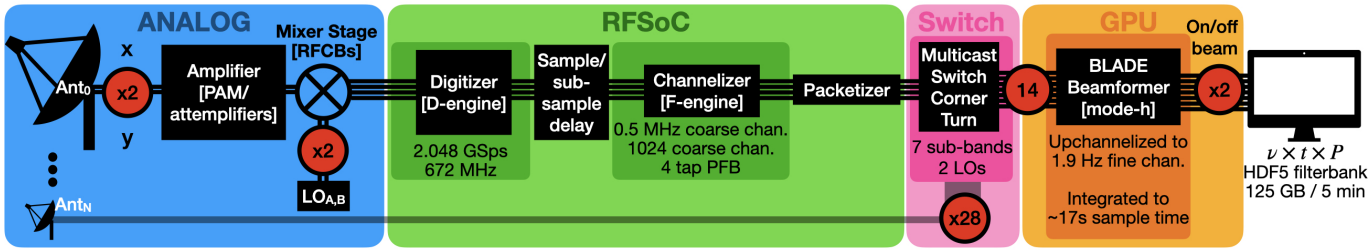


Figure 2. The analog and digital signal path of the ATA as used for this work.

geometry of Figure 1). Thus, even though the monitored regions from the synthesized beams do not cover parts of the ASP to the left and right of the central vertical bisector, all targets in those parts will cross the bisector eventually (or have already crossed it). During an observing session, we loop the 5 minute scans of the ASP (with a variable $\lesssim$60 s downtime between scans) for the entire time that the ASP is above the horizon; this means that we will always have an ecliptic longitude overlap between scans of over 80%, even at the highest frequencies in the survey. We note that this provides far more than 5 minutes of coverage for each actual sky location and that any true signals that we see could show up in multiple contiguous pointings (although this would depend on the duration of the transmission, so we do not rely on it in our sensitivity calculations).

We can also assume that any ETA using the ASP technique would prioritize their transmission timing to align at midtransit (i.e., on the bisector). This implies that the total "coverage" of the ASP in this survey will eventually be 100% (or 360°) in ecliptic longitude and 10% in ecliptic latitude, as 10% is the ratio of the *diameter* of the ASP to the (sum of the) synthesized beams. We aim to improve this coverage in future epochs—see Section 4 for a technical review of available options.

2.2. Analysis Strategy and Event Filtering

The full block diagram of the analog path and digital signal processing chain for this project is shown in Figure 2. For each 5 minute scan of ATA-ASP data, the system will output 28 individual `filterbank` HDF5 files: each of the two tunings is divided into seven frequency chunks, and there are two beam files (on and off) in each chunk. One of these files is 4.5 GB in size. With an average of about 80 scans per night, over 365 nights, this project would generate ∼3.7 PB of total data. This is approximately the storage capacity of the entire HCRO server; it is clearly infeasible to store and process all of these data in a fully offline fashion, and thus a pseudo-real-time processing strategy must be implemented.

The general processing pipeline is as follows: narrowband "hits" (Hz-scale signals with drift rate and signal-to-noise ratio, SNR, values within the search range) will be found by the `bliss` code.[15] These hits will be filtered with an RFI mask and drift rate threshold. Then, the remaining hits will be cross-referenced between the two beams using the `NBeamAnalysis`[16] "event" finding code (N. Tusay et al. 2024), which searches for sky-localized hits that are significantly brighter in one beam than another. Any remaining events after the `NBeamAnalysis` filtering will be inspected visually. If no sky-localized narrowband signals are found in a session, the original `filterbank` data for that night will be deleted, while the reduced data products—.csv files of hits and events, logs, and output plots—will be saved permanently. Conversely, if a signal of interest (a potentially sky-localized narrowband signal) is found in a session, both the original data

[15] The software is available on the GitHub `bliss` codebase, https://github.com/UCBerkeleySETI/bliss, under a 3-Clause BSD license. The version used for this analysis is archived in Zenodo (N. West 2025).

[16] The software is available on the GitHub `NBeamAnalysis` codebase, https://github.com/isabelgerrard/NBeamAnalysis, under a 3-Clause BSD license, and version 2.0 is archived in Zenodo (N. Tusay & I. Gerrard 2026).

and the output products described above will be saved permanently, for the entire session across all frequencies, to HCRO storage. From the pilot observations and other preliminary tests, we find that signals of interest are rare enough that this procedure ensures that the data rate remains manageable while keeping full spectral and temporal context around interesting hits.

2.2.1. bliss

In general, recent narrowband technosignature searches have used either the turboSETI code (E. Enriquez & D. Price 2019), developed by BL, or the UCLA SETI pipeline (J.-L. Margot et al. 2023). On the ATA, only turboSETI had previously been implemented. turboSETI estimates drift rates by "dedrifting" a signal according to a set of trial drift rates and choosing the one that maximizes the dedrifted signal's amplitude when the dedrifted spectrogram is summed over time. The method is computationally intensive but is aided by a fast tree algorithm originally developed by J. H. Taylor (1974) for pulsar dedispersion. A GPU-accelerated version of the turboSETI code, called seticore, is available,[17] and further details of the method are given in J. E. Enriquez et al. (2017).

bliss performs a direct dedrifting routine, i.e., not using the tree algorithm of J. H. Taylor (1974). This results in greater computational complexity but mitigates the dedrifting loss characterized by V. Gajjar et al. (2021); see Section 2.3 for more details. After dedrifting the spectrogram at each trial drift rate, bliss sums over time to obtain a one-dimensional spectrum. All dedrifted spectra are stored sequentially in a drift–frequency plane, analogous to the dispersion measure–time "butterfly plot" plane used in some searches for pulsars and FRBs (e.g., G. Foster et al. 2018). bliss executes its hit search in this plane, which offers several advantages in signal parameter estimation.

It has been recently appreciated that turboSETI has some drawbacks. C. Choza et al. (2023) found that the rolloff from standard coarse channelization causes turboSETI's noise statistics to be skewed, returning a lower estimated mean and higher standard deviation than it should with proper rolloff mitigation. Signals can also be attenuated by tens of percent if they fall near the edge of a fine channel (instead of at its center; see discussion in Section 2.3) and can be attenuated by factors of a few to several without additional mitigation for Doppler smearing at drift rates beyond the "unit drift rate." Overall, C. Choza et al. (2023) discovered with injection-recovery testing that a user-input SNR threshold of 10 actually corresponds to an SNR threshold of 33 times the true noise in the dynamic spectrum in a Green Bank Telescope (GBT) observation.

bliss accounts for signal attenuation due to both coarse-channel rolloff and Doppler smearing at high drift. The shape of the rolloff is deterministic, and, given certain parameters of the signal processing pipeline, its mitigation is as straightforward as modeling the shape and dividing it out of the data in the frequency domain. The bliss pipeline performs this correction as a standard part of its hit search. The correction does not take into account the shape of the bandpass, which can still bias the noise statistics and result in false positives even after the rolloff correction. However, this typically only occurs in rolloff regions and is readily taken care of by our RFI masking procedures. Meanwhile, the direct dedrifting approach taken by bliss permits better characterization of Doppler smearing than the tree algorithm, with the result that much of the signal power lost can be recovered, or "desmeared." A forthcoming study by B. Jacobson-Bell & S. Croft et al. (2026) conducts signal injection and recovery tests using setigen (B. Brzycki et al. 2022) to simulate narrowband signals at various widths, drift rates, and SNRs. The authors find that, while turboSETI consistently underestimates the SNR by a factor of >3.3, and this factor is worse for higher-drift signals, bliss's SNR estimation is ~ 1 for low-drift signals and still an improvement on turboSETI's at high drifts. Injection-recovery testing through the entire ATA pipeline, including NBeamAnalysis, is currently underway and will be reported in future work (I. Gerrard et al. 2026, in preparation).

In addition to the SNR underestimation reported by C. Choza et al. (2023), turboSETI has been criticized for its "top hit search," a routine designed to avoid redundant detections that, in practice, results in the discarding of $\sim$kHz-wide spectral windows surrounding high-SNR signals (P. Pinchuk et al. 2019; J.-L. Margot et al. 2021). This is a very small fraction of the $\sim$GHz-wide bandwidths typical of SETI surveys, but it nevertheless biases our constraints. By leveraging the full drift–frequency plane at once, bliss can disambiguate nearby signals more effectively. A user-selectable parameter designates the minimum L_1 distance between two peaks in this plane for them to be considered separate signals; that is, bliss requires that signals not only align in frequency but also have similar drift rates to be counted as redundant. A search with this parameter is referred to as a "connected components" search. Overlapping signals, then, will typically be missed by turboSETI but found by bliss so long as they have different drift rates.

Another assumption made by turboSETI is that any signals it finds have bandwidths less than the survey's frequency resolution. When turboSETI detects a signal that spans more than one frequency bin, it does not integrate the signal's power across its frequency profile, giving a per-frequency power value that is assumed to be the total power in our calculation of transmitter limits (see Section 2.3). (Indeed, turboSETI has no way of estimating a multibin signal's bandwidth at all!) In contrast, bliss takes into account that multibin signals are possible and estimates each signal's bandwidth and total power accordingly. B. Jacobson-Bell & S. Croft et al. (2026) conduct their injection-recovery tests using simulated signals up to 10 bins wide and find that bliss consistently recovers signals of these bandwidths. In our data, wider signals are possible but typically exhibit substructure in their spectral profiles such that bliss is likely to detect them as multiple narrow (<10 bins) subcomponents.

For the above reasons, we use bliss instead of turboSETI as our narrowband hit-finding code for this search. The direct dedrifting approach contributes to substantially longer runtimes, but our implementation at the ATA still runs faster than the rate of data acquisition.

We are running the postprocessing pipeline with an SNR threshold value of 15, which will stay constant throughout the study duration unless the RFI environment gets significantly worse. It should be noted that the coarse-channel shape of the

[17] The software is available on the GitHub seticore codebase, https://github.com/lacker/seticore, under an MIT license.

ATA is more extreme than that of the GBT, which would lead to an even greater difference between the turboSETI and "actual" SNR, compared to the factor of 3.3 for GBT described above. Therefore, while a bliss SNR threshold of 15 seems higher than some turboSETI searches in the past (e.g., S. Z. Sheikh et al. 2023), it is likely actually *more* sensitive in absolute terms than these comparable searches.

In terms of drift rate, bliss also requires a maximum drift rate to use for narrowband hit searching. The maximum drift rate contribution from Earth's rotational motion (which dominates the orbital component at Earth's orbital distance) is 0.11 nHz[18] (S. Z. Sheikh et al. 2019). For surveys that are performing a search of the sky without restricting to known planetary systems, M. G. Li et al. (2023) recommend a maximum drift rate limit of 0.44 nHz, assuming no correction on the part of the transmitter. If we sum these values, we get a maximum drift rate threshold of 0.55 nHz, or $\sim$2 Hz s^{-1} at the highest frequency in our survey. To add some extra buffer on these assumptions, we will double that value to finalize a maximum drift rate for the ATA-ASP survey of 4 Hz s^{-1}, with the assumption that the survey could recover bright, non-drift-corrected signals from targets that may be on the ecliptic but not related to the ASP. Note, however, that in the ASP schema, we can assume that any signal that we receive is being sent from a system that is deliberately sending signals toward the Earth. Therefore, it is more likely (when compared to a standard open signal search) that the transmitter would correct its own drift, and thus signals of interest at lower drift rates within the range may be particularly interesting.

2.2.2. Analysis Pipeline

Given the considerations in this section, we run the hit-finding code bliss for ATA-ASP with an SNR of 15 and a drift rate range of $\pm$4 Hz s^{-1}, we use coarse-channel shape correction and desmearing correction, and we use the connected components method with an L_1 parameter of seven drift–frequency pixels. The outputs from this operation, after being looped through each 5 minute scan, frequency chunk, and beam, will be saved to .csv files, with each row containing information on a particular hit. All hit .csv files will be saved for the duration of the survey, regardless of whether they have a signal of interest in them—this will allow us to assess RFI and hit statistics throughout the survey.

At this stage, we expect an average of 47,000 hits per 5 minute scan per beam, given the results from the ATA-ASP observations thus far. A significant number of these hits are due to RFI in densely populated parts of the spectra, generally delineated by Federal Communications Commission (FCC) allocations. We visually inspected dynamic spectra from all bliss hits from the first 5 minute scan in fall 2024, and spot checks thereafter, to assess which regions needed RFI masking, i.e., removing all hits within the designated frequencies. We thus determined the following list of frequency regions, where any hits in the designated ranges would be removed from consideration:

1. 2355.75–2359.5 MHz,
2. 2484–2488 MHz,
3. 2497.25–2498.5 MHz,
4. 2590–2609 MHz,
5. 2565.75–2585 MHz,
6. 2705–2800 MHz,
7. 3264–3274 MHz,
8. 3341–3343 MHz,
9. 3492.75–3492.85 MHz, and
10. 3600–3700 MHz.

In total, this covers ranges allocated to wireless communication services and geostationary satellite downlinks, among others, and removes 69.46% of the hits associated with RFI while only masking 18.94% of the band.

After the frequency masking is completed, we input the masked hit .csv files into NBeamAnalysis. NBeamAnalysis performs a series of cross-checks between the on- and off- beams, looking for hits that are sky-localized (N. Tusay et al. 2024). NBeamAnalysis first compares the hit .csv for a given scan, searching for hit entries at near-identical frequencies and drift rates. If one is found, the beam ratio is checked: for a 28-dish array, the ratio between the power in a coherent beam centered on an emitting target compared to the power in the incoherent beam should be about a factor of $\sqrt{N_{\rm ants}} = \sqrt{28} = 5.29$. We can use this as a "worst-case scenario" value for comparing two coherent beams: for any true sky-localized signal, given a pair of beams that are placed at least a few synthesized beamwidths away from each other, we would expect a beam ratio of at least 5.29. In reality, this ratio would likely be much greater, and the exact value will depend on where the off-beam falls within the beam pattern of the ATA. Additionally, there will be cases in which an on-beam detection will be above the SNR threshold but the off-beam contains the signal *below* the SNR threshold, which will not be contained in the hit file (and therefore not identified as RFI in this methodology). To identify these situations, further reduce false positives, and rank the level of interest of events at the end of the pipeline, NBeamAnalysis contains another dynamic-spectrum-based comparison step. The code grabs a snippet of the dynamic spectrum from both beams, centered on the hit in the on-beam. Then, it calculates the noise in each snippet and the SNR of a path corresponding to the drift rate and frequency location of the on-beam hit to obtain the beam ratio from the data directly. It also calculates the dot product of the two snippets, which is reported as the "DOT score" between the two beams. In theory, a bright sky-localized signal should make the two snippets very different and produce a DOT score close to 0, while two identical snippets (more likely for local RFI) would produce a DOT score close to 1.

For the ATA-ASP survey, we run the NBeamAnalysis pipeline on a per-night basis, instead of a full-survey basis, and we modify the module slightly, as we need to process several hours of data from a single observing night as close to real time as possible. We implemented multithreaded disk buffer caching and addressed redundant metadata loading to reduce I/O overhead, integrated JIT (Numba) runtime compilation to mitigate computation overhead, and optimized underutilized CPU parallelization. These changes result in a 92.3% reduction in runtime; i.e., even for the longest nights of the year, the modified NBeamAnalysis pipeline analyzes and plots final hits within 8 hr. We ask NBeamAnalysis to plot only hits with beam ratio values of at least 5.29 and at most 1000 plots per night to be visually inspected, limited by the 1000 lowest DOT scores. All generated plots of top hits and .csv files of all events from NBeamAnalysis will be saved for the entire survey to keep up with the incoming data rate.

[18] Here 1 nHz is equivalent to 1 Hz s^{-1} of drift at 1 GHz. This is a useful unit, as the drift rate imparted by a Doppler acceleration is frequency-dependent.

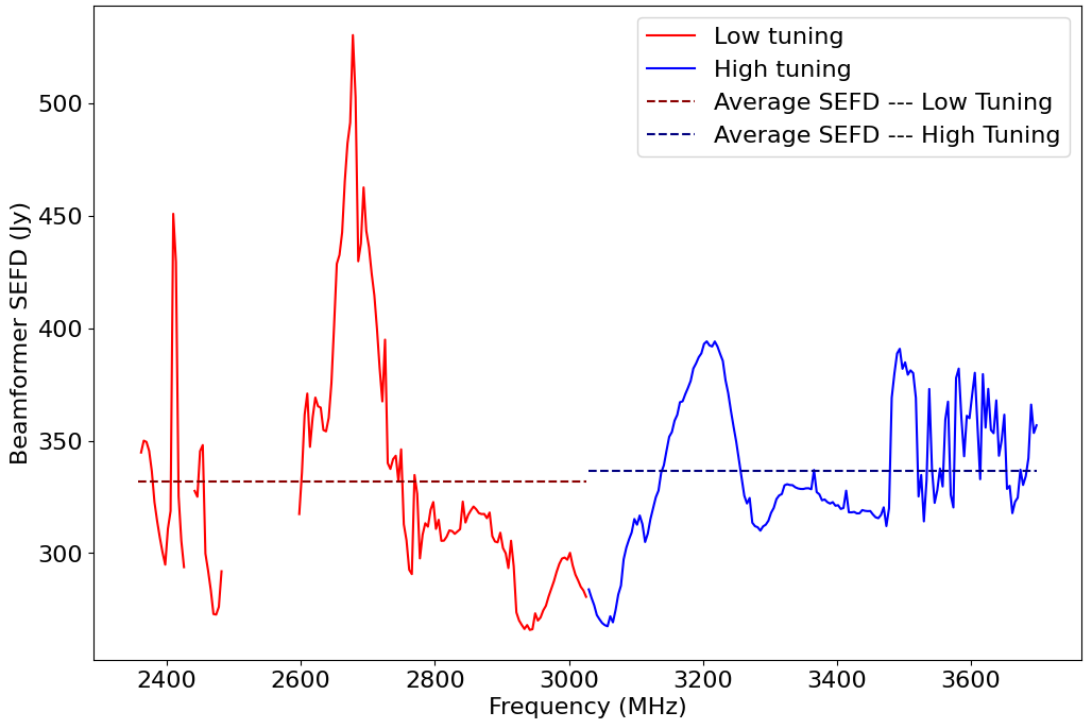


Figure 3. The beamformer SEFD in Jy over frequency in MHz for an example night from the ASP survey. The average SEFD is stable across both halves of the bandpass, but large fluctuations are seen in some regions (e.g., ∼2700 MHz) due to RFI. Missing data from ∼2500–2600 MHz is also due to RFI removal in CASA.

In the case where a potential signal of interest is found, i.e., an event that appears sky-localized and thus warrants follow-up, the entire night of data will be saved. Subsequent follow-up in parallel (following procedures such as those described in, e.g., S. Z. Sheikh et al. 2021; V. Meadows et al. 2022) will occur—one likely situation is a signal of interest appearing sky-localized in one night of data and then reappearing in another night at a different position (thus proving it to be nonsidereal) or being attributed to a particular RFI source. In these cases, the data from the former signal of interest will remain saved on disk for illustrative purposes, but the rest of the data from the night will be deleted.

2.3. Expected Upper Limits and Haystack Fraction

To evaluate the sensitivity of a SETI survey with an interferometer, we must determine the sensitivity of each element in the array and then combine them as appropriate for the beamforming mode. Currently, the ATA normalizes the voltage noise levels to the same Gaussian distribution prior to beamforming and assigns uniform weights $(1 + 0j)$ to all dishes. In the case of similar (but uncorrelated) noise samples per antenna, the beamformer system effective flux density (SEFD) for polarization p is

$$\mathrm{SEFD}_{\mathrm{beamformer},p} = \frac{N}{\left(\sum_{i=1}^{\mathrm{N}} \mathrm{SEFD}_{p,i}^{-0.5}\right)^2}, \tag{1}$$

where the x and y polarizations are thus calculated separately with the equation above and then combined into a total beamformer SEFD via

$$\mathrm{SEFD}_{\mathrm{beamformer,total}} = \frac{1}{\sqrt{\frac{1}{\mathrm{SEFD}_x^2} + \frac{1}{\mathrm{SEFD}_y^2}}}. \tag{2}$$

We use the calibration from an example ASP night (MJD 60893) as a stand-in for the SEFD of the full survey, averaging 334 Jy across the entire bandpass. This value does vary across frequency as shown in Figure 3, but the fluctuations are dominated by RFI rather than an inherent change in sensitivity from the receiver.

If no signals of interest pass verification checks, we will still learn a lot from a nondetection by setting strict new upper limits on radio transmitters using the ASP Schelling point, as well as new limits for the ETZ in general. The minimum detectable flux density of a narrowband radio technosignature search is modified from V. Gajjar et al. (2021) as

$$S_{\mathrm{min,narrow}} = \eta_{\mathrm{DFT}}\eta_{\mathrm{primary}}\frac{\mathrm{SNR}_{\mathrm{min}}}{\sqrt{\beta}} \frac{\mathrm{SEFD}}{\Delta\nu_t}\sqrt{\frac{\Delta\nu}{n_{\mathrm{pol}} \times t_{\mathrm{obs}}}}, \tag{3}$$

where the $\mathrm{SNR}_{\mathrm{threshold}}$ for the full ATA-ASP survey is 15, the SEFD is 334 Jy (as used above), the channel width $\Delta\nu$ is 1.9 Hz, the number of polarizations n_{pol} that the search is sensitive to is 2, the observing time per scan is 5 minutes, and we assume a fiducial transmitter signal width of $\Delta\nu_{t,\mathrm{fid}} = 1$ Hz. We can also calculate the minimum detectable flux $F_{\mathrm{min,narrow}}$ in, e.g., W m^{-2} by simply multiplying both sides by a transmitter signal width of $\Delta\nu_t$. Before plugging values into

this equation, we must consider $\eta_{\rm DFT}$, $\eta_{\rm primary}$, and β in more detail.

The fine channelization is performed via a discrete Fourier transform (DFT), which leads to a sinc function channel response as opposed to an ideal boxcar. Due to this, any signal that is significantly larger than the channel width within one time integration (either intrinsically or due to frequency drift) will recover all of the signal and have a DFT efficiency value $\eta_{\rm DFT} = 1$. Conversely, an infinitely narrowband, nondrifting signal will have an average $\eta_{\rm DFT}$ of 0.77 (the integral of the sinc function over the channel width, as used in E. J. Korpela et al. 2025) and a median $\eta_{\rm DFT}$ of 0.81. In this work, we are looking for narrowband drifting signals with high enough drift rates that we will take $\eta_{\rm DFT} = 1$ for our most likely (and most favorable) case and $\eta_{\rm DFT} = 0.8$ for the least favorable case. Note that there is also some maximum detectable flux above which the fine channel will saturate, but we do not address that in this work.

$\eta_{\rm primary}$ accounts for the response of the primary beam at the location of the synthesized beam where the signal is detected. The on-beam is located at the phase center such that $\eta_{\rm primary} = 1$. But for the off-beam, assuming a Gaussian beam shape, any signals in the off-beam at the lowest frequencies will have 1% attenuation, and those at the highest frequencies will have 2.5% attenuation; in the least favorable case, we thus use $\eta_{\rm primary} = 0.975$.

β is a parameter representing the "dechirping efficiency," which accounts for drift rate smearing past the unit drift rate. For `turboSETI` (the example in V. Gajjar et al. 2021), $S_{\rm min,narrow}$ scales linearly with β, but `bliss` regains some extra signal by including neighboring bins containing smeared signal in the power estimate, reducing the factor to scaling with the square root of β (to first order). The functional form of the dedrifting efficiency is $\beta_{\rm bliss} = \left[\text{round}, (, \frac{\Delta t}{\Delta \nu}, \times, |, \dot{\nu}, |,)\right]^{-1}$. Here, $\dot{\nu}$ is the drift rate being searched, which is rounded to the nearest integer multiple of the unit drift rate $\Delta\nu/\Delta t$.

We make the assumption that ASP beacons are likely drift-compensated (see Section 2.2) on the transmitter's side, but that drift contribution from the Earth's rotation and orbital motion will still be present. As calculated in S. Z. Sheikh et al. (2019), these contributions are at most (with maximally aligned observational geometry) 0.11 nHz. This leads to β values from 1 (when the geometry is favorable) to 0.25 (a reduction of 50% once the square root is taken for Equation (3)) for the highest frequency in the survey.

Plugging these values into Equation (3), we find that a null result over the entire campaign would rule out ASP transmitters with minimum detectable flux densities of approximately 282–1709 in the "best-case" and "worst-case" scenarios. The best case occurs when a signal is located at the center of the primary beam and has a drift rate close to the threshold that narrowly avoids fine-channel attenuation, while the worst case occurs when the signal is located at the off-beam (with primary beam attenuation) and has the highest drift rate in the survey. For context, a transmission dedrifted to the barycenter would have a minimum detectable flux density of 578 Jy even if located at the off-beam. The most powerful transmitter on Earth, the now-defunct planetary radar from the ∼300 m Arecibo telescope, had an effective isotropic radiated power (EIRP) of ∼20 TW; we can thus place limits on Arecibo-scale transmitters with $\Delta\nu_t = 1$ Hz out to 32–79 lightyears in this survey.

ATA-ASP will set the first limits of any kind on the ASP and the most robust continuous transmitter upper limit on the ETZ (covering 10% of its total sky area), which complements the more sensitive, higher-frequency limits on 20 ETZ targets from S. Z. Sheikh et al. (2020).

We can also use the methodology from J. T. Wright et al. (2018) to determine the fraction of the "cosmic haystack" that we will cover in the full ATA-ASP survey, given the parameters described in the previous section. We use provided Python calculation scripts from J. T. Wright et al. (2018) to calculate the haystack fraction using the default haystack bounds and search bounds as follows.

1. A time per scan of 300 s.
2. A minimum frequency of 2.3 GHz.
3. A maximum frequency of 3.7 GHz.
4. A minimum detectable flux density of 282–1709 Jy for the fiducial transmitter signal width of $\Delta\nu_t = 1$ Hz (578 Jy at worst if dedrifted to the barycenter).
5. A channel bandwidth of 1.9 Hz.
6. A polarization fraction of 1 (indicating instrument sensitivity to both polarizations).
7. A sky coverage of 6.8 deg^2. The sky coverage value is obtained by multiplying the width of the highest-frequency (i.e., smallest) synthesized beam size by the 360° circumference of the ecliptic. Then, another factor of 2 is multiplied in to represent the second, also searched, off-beam.

From these input parameters, we find a total haystack volume of 5.1×10^{-21}–3.1×10^{-20}. This volume is about equivalent to the haystack fraction calculated in the previous largest SETI survey ever done with the ATA by G. Harp et al. (2016) and is the same order of magnitude as other large modern search campaigns in similar frequency regions (e.g., R. H. Gray & K. Mooley 2017). Recall as well that even a single additional beam or tuning will increase the sky coverage by 50%, providing ample opportunity for growth of the haystack fraction over the duration of the survey.

3. Pilot Observations

To prepare for the full-scale ASP, we designed a pathfinder survey of a few nights of data, which we ran in summer 2023, with the 20-dish beamformer. In the end, certain decisions that were made for this pilot project (e.g., the frequency tunings) were changed for the final survey based on the results of the pilot. Therefore, we describe the ASP setup as it was performed in summer 2023 in Section 3.1 uniquely from the overall survey design in Section 2.

3.1. Observing Parameters and Data Analysis

For the summer 2024 pilot program, we set the central frequencies of the two frequency tunings to 1.4 and 2.042 GHz, which provided continuous coverage from 1.064 to 2.378 GHz. We selected this range to coincide with the spin-flip frequency of hydrogen at 1420 MHz, which has been proposed as a magic frequency in the past (in fact, as one of the earliest SETI Schelling points; G. Cocconi & P. Morrison 1959). The chosen integration time and frequency resolution for this survey were 15 s and 1 Hz, respectively. The data were taken for an observation time of 5 minutes scan^{-1} with scans looped continuously over the period of time

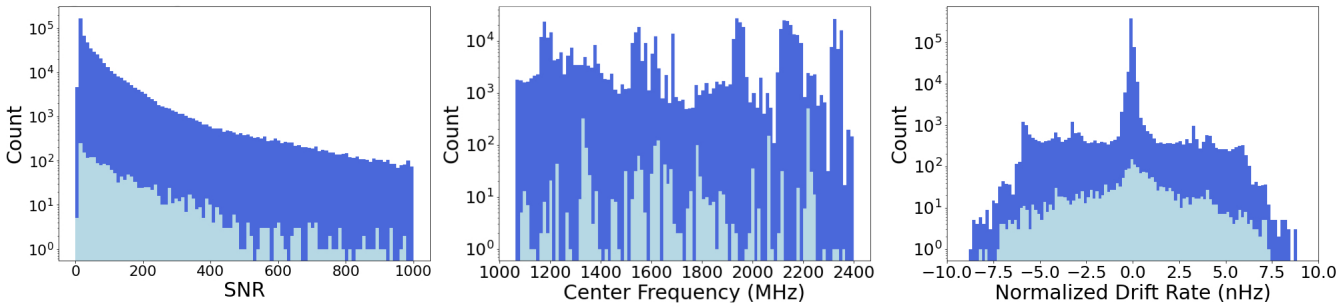


Figure 4. The characteristics of all hits in the ASP pilot survey (dark blue) and the events that were visually inspected (light blue). Generally, the distributions of visually inspected events trace the overall distribution in each variable. Left: histogram of SNR over the course of our 3 nights of observation. The lowest SNR values for visually inspected events on each day range from 10.00 to 10.25, indicating that the spatial filtering maintains the sensitivity from the initial minimum SNR threshold of 10. Center: histogram of center frequencies over the course of our 3 nights of observation. Certain frequency ranges, e.g., 1900–2000 MHz, are much more interference-heavy than others, as expected from frequency allocations in the *L* band. In addition, certain interference-heavy frequency ranges are more accurately removed by the `NBeamAnalysis` spatial filter, leading to the visually inspected events not perfectly tracing the overall structure of the hit distribution. Right: log-scaled histogram of normalized drift rates in nHz for all hits over the course of our 3 nights of observation. As expected, the majority of the hits are centered at zero drift, indicating terrestrial RFI sources.

Table 1
General Search Parameters for the Planned ATA-ASP Campaign and the Pilot Survey

Parameter	Value in Full ATA-ASP Survey	Value in Pilot
	Survey	
Number of dishes	28–42	20
Number of nights	365	3
	Data Product	
Frequencies	2356–3700 MHz	1064–2378 MHz
Total bandwidth	1400 MHz	1400 MHz
Fine-channel width	1.9 Hz	1 Hz
Coarse-channel width	0.5 MHz	0.5 MHz
Time per scan	300 s	300 s
Time integrations	16.8 s	15 s
	Signal Search	
Search software	`bliss`	`turboSETI`
SNR thresholds	15	10
Drift rate range	± 4 Hz s^{-1}	± 10 Hz s^{-1}
Signal bandwidths	∼2–19 Hz	∼1–9 Hz
SEFD	272–291 Jy	332–336 Jy
Minimum detectable flux density	282–1709 Jy	115–936 Jy
Minimum detectable flux at barycenter drift rate (W m^{-2})	3.9×10^{-23}	1.2×10^{-23}
Maximum distance to detect a $\Delta\nu_t = 1$ Hz Arecibo equivalent	32–79 lightyear	44–124 lightyear

Note. The minimum detectable flux at barycenter drift rate (W m^{-2}) is calculated for a signal with a 0.11 nHz normalized drift rate (the maximum from Earth's contribution for an otherwise drift-corrected signal) corrected to the highest frequency in each column, resulting in a signal that smears though a number of channels equivalent to (i.e., has an effective width of) $\Delta\nu_t = 6.8$ Hz for the full campaign and $\Delta\nu_t = 4$ Hz for the pilot survey. As the detected drift rate will often be less than the maximum in the barycentered case, we instead calculate the Arecibo-equivalent distance with $\Delta\nu_t = 1$ Hz throughout this work.

the ASP is visible at the ATA (around 10 p.m.–4 a.m. local time throughout the summer at HCRO). We opted to use two beams as shown in Figure 1, with an on-beam pointed directly at the ASP and an off-beam placed five on-beam widths away from our on-beam as a control.

We collected ∼17 hr of observation of the ASP over the course of 3 nights, which constituted 32 TB of filterbank data. Each file is the same HDF5 "data cube" described in Section 2.1, where each tuning is divided into seven distinct frequency data files. To identify potential narrowband, drifting, sky-localized ASP beacons in these files, we used the `turboSETI` package (E. Enriquez & D. Price 2019) and `NBeamAnalysis` code (as described in more depth in Section 2.2). In this pilot, we selected an SNR threshold of 10 and a maximum drift rate of 10 Hz s^{-1} for `turboSETI`.

General parameters for both the full ATA-ASP survey and the pilot observations in Section 3 are summarized in Table 1.

We then generated plots of the 500 hits with the lowest `NBeamAnalysis` DOT scores on each of the 3 nights, visually inspecting them for signals of interest. Once `NBeamAnalysis` had the beam ratio option available in addition to the DOT score (later in its development), we ran the data through the pipeline again to calculate the beam ratio. We then plotted hits from all 3 nights with a beam ratio above a chosen threshold of 4 (as the pilot program used only 20 dishes, not 28) and visually inspected these 614 hits in addition to the original 1500.

3.2. Results and Lessons Learned

`turboSETI` identified 1.9×10^6 narrowband hits within the 17 hr of data. `NBeamAnalysis` then removed 1.3×10^6 as exact matches between the two beams, leaving us with 5.8×10^5 hits that were initially ranked only on their DOT scores, as this was the only metric implemented in `NBeamAnalysis` at the time. The distributions for normalized drift rate, center frequency, and SNR of both all hits and the visually inspected events for all 3 nights are shown in Figure 4.

Of the remaining 5.8×10^5 hits, we generated dynamic spectra for visual inspection for the 500 hits with the lowest DOT scores on each night. We searched the 1500 produced plots for narrowband and drifting signals that were significantly stronger or exclusively present in the on-beam. The frequency and drift rate characteristics of the visually examined hits are shown in contrast to the entire hit population in Figure 5.

Figure 5 reveals some interesting patterns in how the hits with the lowest DOT scores fall in frequency and drift rate.

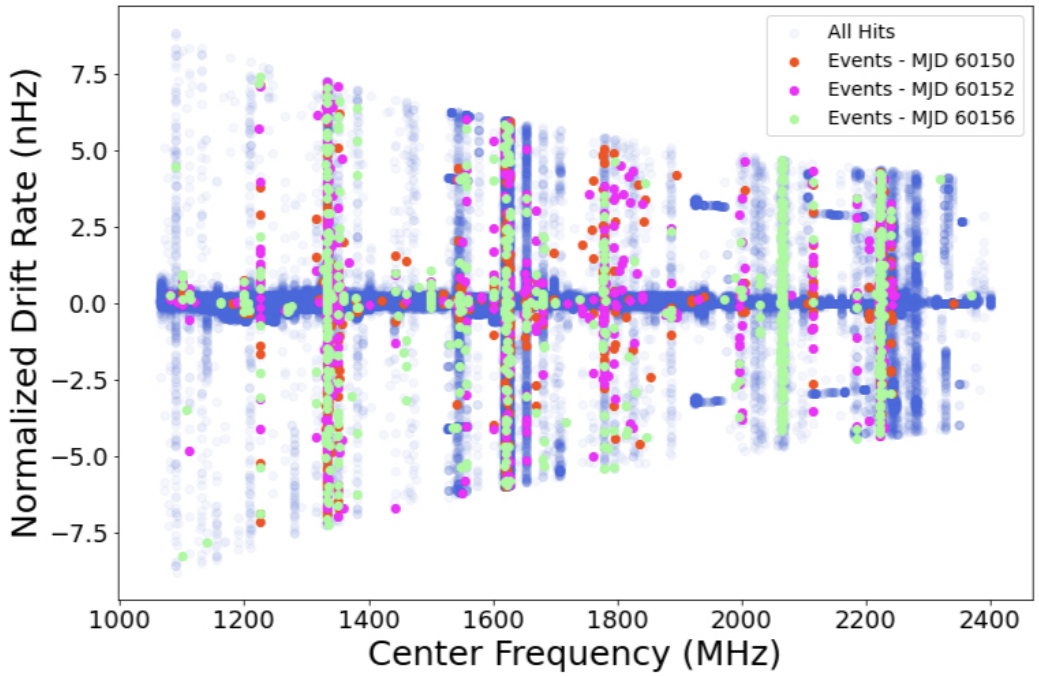


Figure 5. Center frequency plotted against normalized drift rate for all of our hits, as well as the plotted and visually inspected events. Note that (a) most hits fall along a horizontal center line (i.e., their drift rates are near zero); (b) some interference frequency ranges produce hits across all drift rates, not just zero; and (c) the population tapers in drift rate space over the center frequency—this is due to the frequency dependence of input, nonnormalized drift rate versus the normalized plotting. Many plotted events with the lowest DOT scores still fall near a zero drift rate, but some trace high drift interference ranges or, in some cases, appear in relatively empty parts of parameter space.

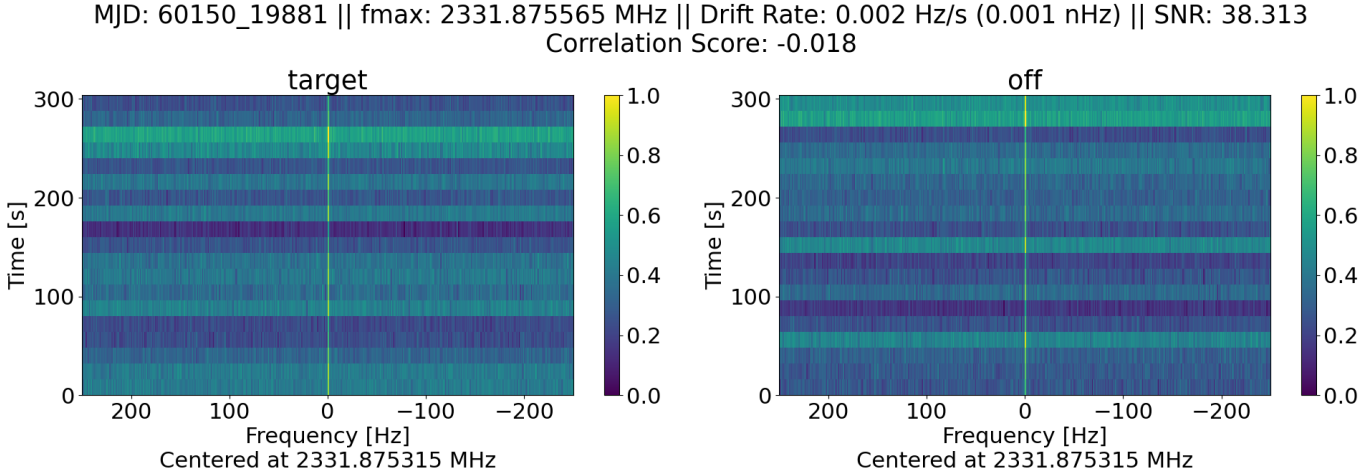


Figure 6. Event that was recorded by `turboSETI` and ranked as potentially sky-localized by `NBeamAnalysis` but was determined to be RFI due to its nearly equal visibility in both beams and minuscule drift rate. This plot is representative of the majority of hits from our observations. It was likely flagged by `NBeamAnalysis` due to the changing broadband background noise level throughout the observation.

Hits with the lowest DOT scores generally do not cluster around a zero drift rate, unlike the overall hit distribution; this is sensible, given that hits at a zero drift rate are likely to be ground-based RFI entering through a side lobe. RFI of this type will be equally bright in the on- and off-beams, thereby warranting a high DOT score. Vertical lines in Figure 5 indicate frequency ranges with RFI that triggers a `turboSETI` hit at every trial drift rate (usually broadband or morphologically complex RFI). These lines are an artifact of the hit-finding process and also appear to a lesser degree in `bliss`. In this reduced view of the data, the most interesting hits are therefore the points in Figure 5 that do not fall along vertical or horizontal features but instead appear relatively isolated in the parameter space. Even so, the full population of highlighted hits in Figure 5 was visually inspected (not just the isolated points).

After visually inspecting our plots, we found no signals that were indicative of a radio technosignature from an ETA. All of the hits produced by our analysis can be explained by terrestrial RFI. The vast majority of signals plotted were narrowband, as expected, but were equally bright in both beams and displayed minimal drift rates, indicative of terrestrial RFI. Figure 6 shows an example output dynamic spectrum from one of these signals.

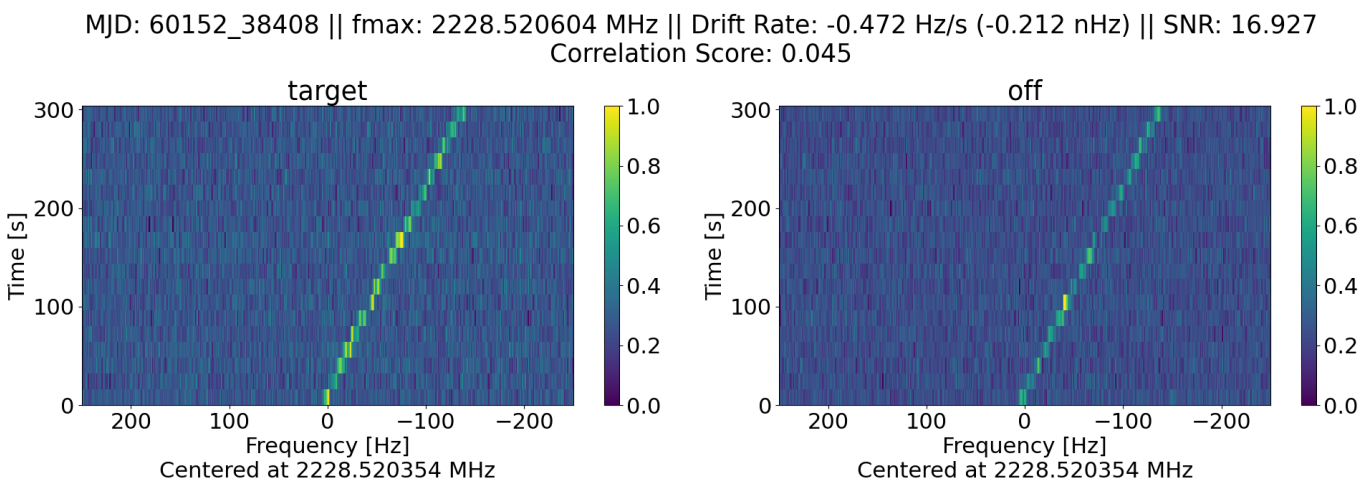


Figure 7. Drifting narrowband signal flagged by `NBeamAnalysis`. While the signal is definitively technological (narrowband) and drifting, its low beam ratio between the two beams and allocation are more consistent with RFI from satellite communications or tracking.

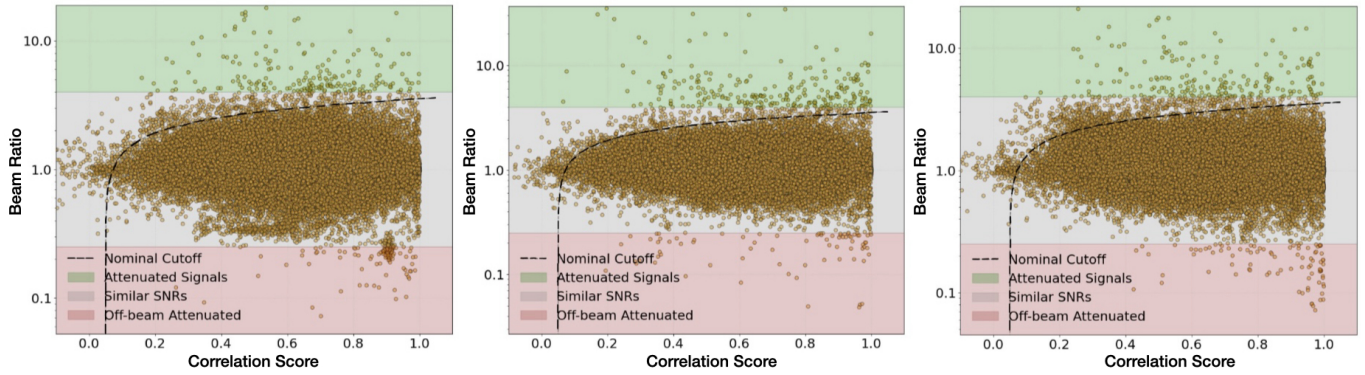


Figure 8. Population-level plots for the 3 nights of data. On the x-axis: DOT score. On the y-axis: beam ratio. The gray and red regions represent hits with beam ratios below the chosen cutoff of 4. The nominal cutoff line is the cutoff for the 500 hits with the lowest DOT score, and the green region contains the hits with beam ratios above the cutoff of 4.

Of the minority of signals that were both drifting and narrowband, none of them showed the expected characteristics of true sky localization. Many of them were sets of aliases that appeared multiple times across the band and were easily cross-referenced with known RFI via the FCC Table of Frequency Allocations. Figure 7 shows one of the "most interesting" signals in our sample, as it is an example of a narrowband drifting signal that is somewhat different in SNR between the on- and off-beams. After further consideration, this signal does not show a large enough beam ratio to be from a distant point source and is most likely RFI within the allocation used for ground-to-space, space-to-ground, and space-to-space transmissions for tracking, telemetry, and communications with satellites.

Figure 8 shows population-level plots for each of the 3 nights, with the dots in the green portion of the plot representing the hits above the beam ratio cutoff of 4 that were subsequently plotted.

Once again, all 614 plots generated in this second round of data analysis fell into the categories of noise or signals characteristic of ground-based RFI, leaving no signals of further interest.

We found no signals of interest in this pilot search of the ASP, a null result that we can convert into an upper limit. The average SEFD of the ATA during our nights of observation was 272 Jy for our 1.4 GHz tuning and 291 Jy for our 2.042 GHz tuning; we take an average value of 281.5 Jy for the following calculation. The minimum detectable flux density of a narrowband radio technosignature search is given in Equation (3). The $\mathrm{SNR}_{\mathrm{threshold}}$ for the pilot survey is 10, the channel width $\Delta\nu$ is 1 Hz, the number of polarizations n_{pol} that the search is sensitive to is 2, and the observing time per scan is 5 minutes, with an assumed transmitter signal width of Δf_t = 1 Hz. This formulation implies minimum detectable flux densities with $\Delta\nu_t = 1$ Hz of 115–936 Jy for the same best- and worst-case scenarios, respectively, as discussed in Section 2.3. This implies that an Arecibo-like transmitter with $\Delta\nu_t = 1$ Hz could be located as far away as 44–124 lyr and still be detectable in this survey.

In conclusion, we took ∼17 hr of data, which represents 0.8% of the full ATA-ASP survey described in Section 2. From those hours of observation, we found 1.9×10^6 total hits and were left with 5.8×10^5 hits after removing identical matches with `NBeamAnalysis`. We visually inspected the top 500 hits from each night and found zero sky-localized signals of interest. From this, we can conclude that we find no narrowband technosignatures in the R.A. range 20.266–21.312 on the ASP above a flux density threshold of 115–936 Jy at these frequencies.

More generally, the pilot survey produced a series of "lessons learned" that influenced the full ATA-ASP survey design in Section 2.

1. *Frequency range.* The *L*-band frequencies in the pilot survey were extremely RFI-dense, which caused the search algorithms to take an order of magnitude longer to run than would be needed for ATA-ASP. For the full survey, we opted for slightly higher frequencies (at the cost of a reduced coverage area) to avoid this issue.
2. *RFI masking.* The pilot survey showed that we could identify certain "problem" allocations within the spectrum that were relatively narrow in bandwidth but would produce the majority of the hits in the survey. This inspired the RFI masking step that is being implemented on the full survey.
3. *Beam ratio.* We found that the DOT score in `NBeamAnalysis`, while useful as a stand-alone metric, is more diagnostic of a signal's nonterrestrial origin if combined with a calculated beam ratio, as discovered in tests with Mars orbiters described in N. Tusay et al. (2024). We are implementing the combined beam ratio plus DOT score filtering method in the full ATA-ASP survey.
4. *Channel width and file size.* Based on the data rate calculations from the pilot, we identified that slightly increasing the channel width and integration size would halve the file size, allowing for more long-term storage options.

4. Conclusions and Future Work

The ASP, a Schelling point in time along the ETZ bracketing the ecliptic, may be a particularly well-motivated spacetime coordinate in the search for radio technosignatures. While the idea of searching for beacons in the ASP has been present in the literature for some time, the actual observational search coverage is essentially zero. In this paper, we described a new campaign with the ATA, titled ATA-ASP, which endeavors to perform the first large-scale search of this region in a multiyear *S*-band observing campaign. We also presented a null result from pilot observations taken over 3 nights in summer 2023, which demonstrated the technology and procedure for performing ATA-ASP and set the first upper limits on Arecibo-like transmitters with $\Delta\nu_t = 1$ Hz out to ∼44–124 lightyears for the ASP. When the full survey is completed, assuming two synthesized beams, we will have covered a haystack fraction of 5.1×10^{-21}–3.1×10^{-20} and placed limits on Arecibo-scale transmitters with $\Delta\nu_t = 1$ Hz out to 32–79 lightyears.

As shown in Figure 1, the two-beam ATA-ASP survey, even when completed, will only cover ∼10% of the ASP in ecliptic latitude, or 20% if limited to the restricted ETZ (R. Heller & R. E. Pudritz 2016). As such, we will endeavor to increase the latitude coverage as additional ATA capabilities come online during the course of the survey. One option for expanding the latitude coverage is to add additional beams along the ASP bisector, as shown in Figure 1. Currently, the corresponding increase in data rate and processing time is infeasible at these frequency resolutions and bandwidths: the frequency resolution would have to be increased by 250× and the bandwidth decreased by 6× to produce the ∼20 beams that would be needed to span the ASP. However, as more compute power is added to the backend machines and the `NBeamAnalysis` code is optimized over the next few years, adding additional beams may be possible. A second option for expanding the latitude coverage is to perform the narrowband search within the "image plane," i.e., to use the in-development BLADE correlator backend on the ATA to create a sky image of the entire ASP. This spectral line correlator is still undergoing testing, and there are no existing open-source algorithms for performing such a search. However, as prototyping of an image-plane narrowband search code begins over the next few years, it could be ported to a future epoch of ATA-ASP.

Thus far, we have taken a total of >121 nights of nonpilot data. Analysis of the first year of ASP data will be presented in an upcoming publication.

Acknowledgments

The Allen Telescope Array (ATA) refurbishment program and its ongoing operations have received substantial support from Franklin Antonio. Additional contributions from Frank Levinson, Greg Papadopoulos, the Breakthrough Listen Initiative, and other private donors have been instrumental in the renewal of the ATA. Breakthrough Listen is managed by the Breakthrough Initiatives, sponsored by the Breakthrough Prize Foundation. The Paul G. Allen Family Foundation provided major support for the design and construction of the ATA, alongside contributions from Nathan Myhrvold, Xilinx Corporation, Sun Microsystems, and other private donors. The ATA has also been supported by contributions from the US Naval Observatory and the US National Science Foundation. S.Z.S. acknowledges that this material is based upon work supported by the National Science Foundation MPS-Ascend Postdoctoral Research Fellowship under grant No. 2138147. E. H.'s participation was enabled through the SETI Institute REU program funded under NSF award No. 2051007.

ORCID iDs

Sofia Z. Sheikh https://orcid.org/0000-0001-7057-4999
Bárbara Cabrales https://orcid.org/0009-0009-1019-3890
Isabel Gerrard https://orcid.org/0009-0003-9610-5068
Wael Farah https://orcid.org/0000-0002-0161-7243
Nick Tusay https://orcid.org/0000-0001-9686-5890
David R. DeBoer https://orcid.org/0000-0003-3197-2294
Saul Balcarcel-Salazar https://orcid.org/0009-0005-1104-3502
Joe S. Bright https://orcid.org/0000-0002-7735-5796
Luigi F. Cruz https://orcid.org/0000-0001-5576-2254
Ben Jacobson-Bell https://orcid.org/0009-0009-6231-9280
Alexander W. Pollak https://orcid.org/0000-0002-3430-7671
Andrew Siemion https://orcid.org/0000-0003-2828-7720

References

Bright, J. S., Rhodes, L., Farah, W., et al. 2023, NatAs, 7, 986
Brzycki, B., Siemion, A. P. V., Pater, I. D., et al. 2022, AJ, 163, 222
Cabrales, B., Davenport, J. R., Sheikh, S., et al. 2024, AJ, 167, 101
Castellano, T., Doyle, L., & McIntosh, D. 2004, in Symposium-International Astronomical Union, Vol. 202 (Cambridge Univ. Press), 445
Choza, C., Bautista, D., Croft, S., et al. 2023, AJ, 167, 10
Cocconi, G., & Morrison, P. 1959, Natur, 184, 844
Colomb, F., Hurrell, E., Lemarchand, G., & Olald, J. 1995, ASPC, 74, 345
Corbet, R. H. 1999, PASP, 111, 881
Corbet, R. H. 2003, AsBio, 3, 305

Cruz, L. 2024, luigifcruz/blade: BLADE: Breakthrough Listen Accelerated DSP Engine, v1.0.5, Zenodo, doi:10.5281/zenodo.21454459
Davenport, J. R., Cabrales, B., Sheikh, S., et al. 2022, AJ, 164, 117
Ekers, R. D., Culler, K., Billingham, J., & Scheffer, L. 2002, SETI 2020: a Roadmap for the Search for Extraterrestrial Intelligence (SETI Press)
Enriquez, E., & Price, D. 2019, turboSETI: Python-based SETI search algorithm, Astrophysics Source Code Library, ascl:1906.006
Enriquez, J. E., Siemion, A., Foster, G., et al. 2017, ApJ, 849, 104
Foster, G., Karastergiou, A., Geyer, M., et al. 2018, MNRAS, 481, 2612
Freitas, R. 1980, JBIS, 33, 95
Gajjar, V., Perez, K. I., Siemion, A. P., et al. 2021, AJ, 162, 33
Gray, R. H., & Mooley, K. 2017, AJ, 153, 110
Harp, G., Richards, J., Tarter, J. C., et al. 2016, AJ, 152, 181
Heller, R., & Pudritz, R. E. 2016, AsBio, 16, 259
Horowitz, P., Coldwell, C. M., Howard, A. B., et al. 2001, SPIE, 4273, 119
Jacobson-Bell, B., Croft, S., et al. 2026, BAAS, 247, 105.06
Kipping, D., Frank, A., & Scharf, C. 2020, IJAsB, 19, 430
Korpela, E. J., Anderson, D. P., Cobb, J., et al. 2025, AJ, 170, 112
Leigh, D., & Horowitz, P. 2000, in A New Era in Bioastronomy, Vol. 213, (ASP), 459
Li, M. G., Sheikh, S. Z., Gilbertson, C., et al. 2023, AJ, 166, 182
Machol, R. E., van Vleck, E., & Evans, L. C. 1975, in ICIASF 1975; 6th International Congress on Instrumentation in Aerospace Simulation Facilities (IEEE), 17
Maire, J., Wright, S. A., Holder, J., et al. 2022, SPIE, 12184, 2686
Makovetskii, P. 1980, Icar, 41, 178
Margot, J.-L., Li, M. G., Pinchuk, P., et al. 2023, AJ, 166, 206
Margot, J.-L., Pinchuk, P., Geil, R., et al. 2021, AJ, 161, 55
McLaughlin, W. I. 1977, Icar, 32, 464
Meadows, V., Graham, H., Abrahamsson, V., et al. 2022, arXiv:2210.14293
Nhan, B. D., De Pree, C. G., Iverson, M., et al. 2024, ApJL, 971, L49
Pinchuk, P., Margot, J.-L., Greenberg, A. H., et al. 2019, AJ, 157, 122
Prusti, T., De Bruijne, J., Brown, A. G., et al. 2016, A&A, 595, A1
Saide, R. C., Garrett, M. A., & Heeralall-Issur, N. 2023, MNRAS, 522, 2393
Schelling, T. C. 1980, The Strategy of Conflict (Harvard Univ. Press)
Schirru, L., Ladu, A., & Gaudiomonte, F. 2023, Univ, 9, 390
Sheikh, S. Z. 2020, IJAsB, 19, 237
Sheikh, S. Z., Farah, W., Pollak, A. W., et al. 2024, MNRAS, 527, 10425
Sheikh, S. Z., Kanodia, S., Lubar, E., et al. 2023, AJ, 165, 61
Sheikh, S. Z., Siemion, A., Enriquez, J. E., et al. 2020, AJ, 160, 29
Sheikh, S. Z., Smith, S., Price, D. C., et al. 2021, NatAs, 5, 1153
Sheikh, S. Z., Wright, J. T., Siemion, A., & Enriquez, J. E. 2019, ApJ, 884, 14
Shostak, S., & Villard, R. 2004, in Symposium-International Astronomical Union, Vol. 213, (Cambridge University Press), 409
Singer, C. E. 1982, CosSe, 4, 22
Tarter, J. 2001, ARA&A, 39, 511
Taylor, J. H. 1974, A&AS, 15, 367
Tusay, N., & Gerrard, I. 2026, isabelgerrard/NBeamAnalysis: Efficient and Stable Sky-Localization, v2.0.0, Zenodo, doi:10.5281/zenodo.19188799
Tusay, N., Sheikh, S. Z., Sneed, E. L., et al. 2024, AJ, 168, 283
Welch, J., Backer, D., Blitz, L., et al. 2009, IEEEP, 97, 1438
Wells, R., Poppenhaeger, K., Watson, C., & Heller, R. 2018, MNRAS, 473, 345
West, N. 2025, UCBerkeleySETI/bliss: Breakthrough Listen Interesting Signal Search, 2025-05-22, Zenodo, doi:10.5281/zenodo.19131179
Worden, S. P., Drew, J., Siemion, A., et al. 2017, AcAau, 139, 98
Wright, J. T. 2020, IJAsB, 19, 446
Wright, J. T., Kanodia, S., & Lubar, E. 2018, AJ, 156, 260
Zuckerman, B. 1985, AcAau, 12, 127