Journal of
Astronomical Telescopes,
Instruments, and Systems

AstronomicalTelescopes.SPIEDigitalLibrary.org

WFIRST-AFTA coronagraph science yield modeling with EXOSIMS

Dmitry Savransky
Daniel Garrett

SPIE.

WFIRST-AFTA coronagraph science yield modeling with EXOSIMS

Dmitry Savransky[a,b,*] and Daniel Garrett[a]
[a]Cornell University, Sibley School of Mechanical and Aerospace Engineering, Ithaca, New York 14853, United States
[b]Cornell University, Carl Sagan Institute, Ithaca, New York 14853, United States

Abstract. We present and discuss the design details of an extensible, modular, open-source software framework called EXOSIMS (Exoplanet Open-Source Imaging Mission Simulator), which creates end-to-end simulations of space-based exoplanet imaging missions. We motivate the development and baseline implementation of the component parts of this software with models of the wide-field infrared survey telescope-astrophysics focused telescope assets (WFIRST-AFTA) coronagraph and present initial results of mission simulations for various iterations of the WFIRST-AFTA coronagraph design. We present and discuss two sets of simulations. The first compares the science yield of completely different instruments in the form of early competing coronagraph designs for WFIRST-AFTA. The second set of simulations evaluates the effects of different operating assumptions, specifically the assumed postprocessing capabilities and telescope vibration levels. We discuss how these results can guide further instrument development and the expected evolution of science yields. © *The Authors. Published by SPIE under a Creative CommonsAttribution 3.0 Unported License. Distribution or reproduction of this work in whole or in part requiresfull attribution of the original publication, including its DOI.* [DOI: 10.1117/1.JATIS.2.1.011006]

Keywords: wide-field infrared survey telescope-astrophysics focused telescope assets; coronography; science yield modeling.

Paper 15030SS received Jun. 1, 2015; accepted for publication Nov. 6, 2015; published online Dec. 10, 2015.

1 Introduction

The majority of exoplanets discovered to date have been detected indirectly, by looking for effects these planets have on their host stars. Directly imaging exoplanets will provide a great deal of additional information unobtainable by most indirect detection methods and make discoveries expanding the population of known exoplanets. While direct imaging of exoplanets has been demonstrated with ground-based instruments, these have all been very young, very large, and self-luminous planets on long-period orbits. Imaging of smaller and more Earth-like planets will likely require space observatories, such as the wide-field infrared survey telescope-astrophysics focused telescope assets (WFIRST-AFTA). Such observatories are major undertakings requiring extensive planning and design.

Building confidence in a mission concept's ability to achieve its science goals is always desirable. Unfortunately, accurately modeling the science yield of an exoplanet imager can be almost as complicated as designing the mission. While each component of the system is modeled in great detail as it proceeds through its design iterations, fitting these models together is very challenging. Making statements about expected science returns over the course of the whole mission requires a large number of often unstated assumptions when such results are presented. This makes it challenging to compare science simulation results and also to systematically test the effects of changing just one part of the mission or instrument design from different groups.

We seek to address this problem with the introduction of a new modular, open-source mission simulation tool called EXOSIMS (Exoplanet Open-Source Imaging Mission Simulator). This software is specifically designed to allow for systematic exploration of exoplanet imaging mission science yields. The software framework makes it simple to change the modeling of just one aspect of the instrument, observatory, or overall mission design. At the same time, this framework allows for rapid prototyping of completely new mission concepts by reusing pieces of previously implemented models from other mission simulations.

Modeling the science yield of an exoplanet imager is primarily difficult because it is completely conditional on the true distributions of planet orbital and physical parameters, of which we so far have only partial estimates. This makes the mission model an inherently probabilistic one, which reports posterior distributions of outcomes conditioned on some selected priors. Since the introduction of observational completeness by Brown,[1] it is common to approach exoplanet mission modeling with Monte Carlo methods. Various groups have pursued such modeling, often focusing on specific aspects of the overall mission or observation modeling.[2–5]

A second challenge is correctly including all of the dynamic and stochastic aspects of such a mission. Given a spacecraft orbit, a target list, and the constraints of the imaging instrument, we can always predict when targets will be observable. Incorporating this knowledge into a simulation, however, can be challenging if a single calculated value represents the predictions, i.e., the number of planets discovered. Similarly, while it is simple to write down the probability of detecting a planet upon the first observation of a star, it is more challenging to do the same for a second observation an arbitrary amount of time later, without resorting to numerical simulation.[2] EXOSIMS deals with these challenges by explicitly simulating every aspect of the mission and producing a complete timeline of simulated observations including the specific targets observed at specific times in the mission and recording the simulated outcomes of these observations. While one such simulation does not answer

*Address all correspondence to: Dmitry Savransky, E-mail: ds264@cornell.edu

the question of expected mission science yield, an ensemble of many thousands of such simulations gives the data for the posterior distributions of science yield metrics. EXOSIMS is designed to generate these ensembles and provide the tools to analyze them, while allowing the user to model any aspect of the mission in as much detail as desired.

In Sec. 2, we provide an overview of the software framework and some details on its component parts. As the software is intended to be highly reconfigurable, we focus on the operational aspects of the code rather than implementation details. We use the coronagraphic instrument currently being developed for WFIRST-AFTA as a motivating example for specific implementations of the code. In Sec. 3, we present mission simulation results for various iterations of the WFIRST-AFTA coronagraph designs using components that are being adapted to build the final implementation of EXOSIMS.

EXOSIMS is currently being developed as part of a WFIRST Preparatory Science investigation, with initial implementation targeted at WFIRST-AFTA. This development includes the definition of a strict interface control, along with corresponding prototypes and class definitions for each of the modules described below. The interface control document and as-built documentation are both available for public review and comment.[6] Initial code release is targeted for fall 2015, with an alpha release in February of 2016 and continued updates through 2017.

Future development of EXOSIMS is intended to be a community-driven project, and all software related to the base module definitions and simulation execution will be made publicly available alongside the interface control documentation to allow mission planners and instrument designers to quickly write their own modules and drop them directly into the code without additional modifications made elsewhere. We fully expect that EXOSIMS will be highly useful for ensuring the achievement of the WFIRST-AFTA science goals and will be of use in the design and planning of future exoplanet imaging missions.

2 EXOSIMS Description

EXOSIMS builds upon previous frameworks described in Refs. 3 and 7, but will be significantly more flexible than these earlier efforts, allowing for seamless integration of independent software modules, each of which performs its own well-defined tasks, into a unified mission simulation. This will allow the wider exoplanet community to quickly test the effects of changing a single set of assumptions (for example, the specific model of planet spectra, or a set of mission operating rules) on the overall science yield of the mission by only updating one part of the simulation code rather than rewriting the entire simulation framework.

The terminology used to describe the software implementation is loosely based on the object-oriented framework upon which EXOSIMS is built. The term module can refer to either the object class prototype representing the abstracted functionality of one piece of the software, or to an implementation of this object class, which inherits the attributes of the prototype, or to an instance of this object class. Thus, when we speak of input/output definitions of modules, we are referring to the class prototype. When we discuss implemented modules, we mean the inherited class definition. Finally, when we speak of passing modules (or their outputs), we mean the instantiation of the inherited object class being used in a given simulation. Relying on strict inheritance for all implemented module classes provides an automated error and consistency-checking mechanism, as we can always compare the outputs of a given object instance to the outputs of the prototype. This means that it is trivial to precheck whether a given module implementation will work with the larger framework, and thus allows for the flexibility and adaptability described above.

Figure 1 shows the relationships of the component software modules classified as either input modules or simulation modules. The input modules contain specific mission design parameters. The simulation modules take the information contained in the input modules and perform mission simulation tasks. Any module may perform any number or kind of calculations using any or all of the input parameters provided. They are only constrained by their input and output specifications, which are designed to be as flexible as possible, while limiting unnecessary data passing to speed up execution.

2.1 *Input Modules*

The specific mission design under investigation determines the functionality of each of the input modules, but the inputs and outputs of each are always the same (in terms of data type and what the variables represent). These modules encode and/or generate all of the information necessary to perform mission simulations. Here we briefly describe the functionality and major tasks for each of the input module.

2.1.1 *Optical system description*

The optical system description module contains all of the necessary information to describe the effects of the telescope and starlight suppression system on the target star and planet wavefronts. This requires encoding the design of both the telescope optics and the specific starlight suppression system, whether it be an internal coronagraph or an external occulter. The encoding can be achieved by specifying point spread functions (PSFs) for on- and off-axis sources, along with (potentially angular separation-dependent) contrast and throughput definitions. At the opposite level of complexity, the encoded portions of this module may be a description of all of the optical elements between the telescope aperture and the imaging detector, along with a method of propagating an input wavefront to the final image plane. Intermediate implementations can include partial propagations, or collections of static PSFs representing the contributions of various system elements. The encoding of the optical train will allow for the extraction of specific bulk parameters, including the instrument inner working angle (IWA), outer working angle (OWA), and mean and max contrast and throughput.

If the starlight suppression system includes active wavefront control, i.e., via one or more deformable mirrors (DM),[8] then this module must also encode information about the sensing and control mechanisms. Again, this can be achieved by simply encoding a static targeted DM shape or by dynamically calculating DM settings for specific targets via simulated phase retrieval. As wavefront control residuals may be a significant source of error in the final contrast budget, it is vitally important to include the effects of this part of the optical train.

The optical system description can optionally include stochastic and systematic wavefront-error generating components. Again, there is a wide range of possible encodings and complexities. They could be Gaussian errors on the contrast curves

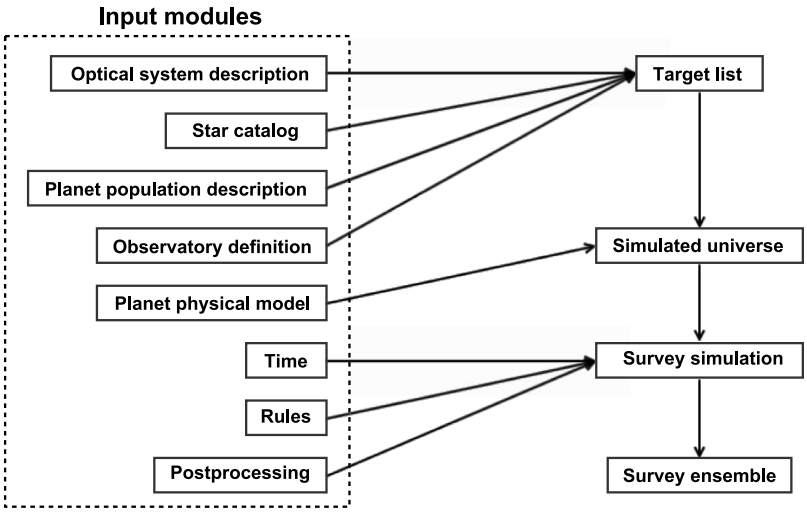

Fig. 1 Flow chart of mission simulation. Each box represents a component software module that interacts with other modules as indicated by the arrows. The simulation modules (those that are not classified as input modules) pass all input modules along with their own output. Thus, the survey ensemble module has access to all of the input modules and all of the upstream simulation modules.

sampled during survey simulation to add a random element to the achieved contrast on each target. Alternatively, in cases where an active wavefront control system is modeled, stochastic wavefront errors could be introduced by simulating the measurement noise on the wavefront sensor (either again as drawn from predetermined distributions or additively from various detector and astrophysical noise sources). Systematic errors, such as miscalibration of DM, closed-loop control delays, and noncommon path errors, may be included to investigate their effects on contrast or optical system overhead. In cases where the optical system is represented by collections of static PSFs, these effects must be included in the diffractive modeling that takes place before executing the simulation. For external occulters, we draw on the large body of work on the effects of occulter shape and positioning errors on the achieved contrast, as in Ref. 9.

Finally, the optical system description must also include a description of the science instrument or instruments. The baseline instrument is assumed to be an imaging spectrometer, but pure imagers and spectrometers are also supported. Each instrument encoding must provide its spatial and wavelength coverage and sampling. Detector details such as read noise, dark current, and quantum efficiency must be provided, along with more specific quantities such as clock-induced charge for electron multiplying CCDs.[10] Optionally, this portion of the module may include descriptions of specific readout modes, i.e., in cases where Fowler sampling[11] or other noise-reducing techniques are employed. In cases where multiple science instruments are defined, they are given enumerated indices in the specification, and the survey simulation module must be implemented so that a particular instrument index is used for a specific task, i.e., detection versus characterization.

The overhead time of the optical system must also be provided and is split into two parameters. The first is an integration time multiplier for detection and characterization modes, which represents the individual number of exposures that need to be taken to cover the full field of view, full spectral band, and all polarization states in cases where the instrument splits polarizations. For detection modes, we will typically wish to cover the full field of view, while possibly only covering a small bandpass and only one polarization, whereas for characterizations, we will typically want all polarizations and spectral bands, while focusing on only one part of the field of view. The second overhead parameter gives a value for how long it will take to reach the instrument's designed contrast on a given target. This overhead is separate from the one specified in the observatory definition, which represents the observatory settling time and may be a function of orbital position, whereas the contrast floor overhead may depend on target brightness. If this value is constant, as in the case of an observing strategy where a bright target is used to generate the high-contrast regions, or zero, as in the case of an occulter, then it can be folded in with the observatory overhead.

2.1.2 *Star catalog*

The star catalog module includes detailed information about potential target stars drawn from general databases such as SIMBAD,[12] mission catalogs such as hipparcos,[13] or from existing curated lists specifically designed for exoplanet imaging missions.[4] Information to be stored or accessed by this module will include target positions and proper motions at the reference epoch (see Sec. 2.1.6), catalog identifiers (for later cross-referencing), bolometric luminosities, stellar masses, and magnitudes in standard observing bands. When direct measurements of any value are not available, values are synthesized from ancillary data and empirical relationships, such as color relationships and mass–luminosity relations.[14]

This module will not provide any functionality for picking the specific targets to be observed in any one simulation, nor even for culling targets from the input lists where no observations of a planet could take place. This is done in the target list module as it requires interactions with the planetary population module (to determine the population of interest), the optical system description module (to define the capabilities of the instrument), and observatory definition module (to determine if the view of the target is unobstructed).

2.1.3 *Planet population description*

The planet population description module encodes the density functions of all required planetary parameters, both physical and orbital. These include semimajor axis, eccentricity, orbital orientation, and planetary radius and mass. Certain parameter models may be empirically derived,[15] while others may come from analyses[16,17] of observational surveys, such as the Keck Planet Search,[18,19] Kepler,[20–22] and ground-based imaging surveys including the Gemini Planet Imager Exoplanet Survey.[23,24] This module also encodes the limits on all parameters to be used for sampling the distributions and determining derived cutoff values such as the maximum target distance for a given instrument's IWA.

The planet population description module does not model the physics of planetary orbits or the amount of light reflected or emitted by a given planet, but rather only encodes the statistics of planetary occurrence and properties. As this encoding is based on density functions, it fully supports modeling "toy" universes where all parameters are fixed, in which case all of the distributions become delta functions. We can equally use this encoding to generate simulated universes containing only "Earth-twins" to compare with previous studies as in Ref. 1 or 5. Alternatively, the distributions can be selected to mirror, as closely as possible, the known distributions of planetary parameters. As this knowledge is limited to specific orbital or mass/radius scales, this process invariably involves some extrapolation.

2.1.4 *Observatory description*

The observatory definition module contains all of the information specific to the space-based observatory not included in the optical system description module. The module has three main tasks: orbit, duty cycle, and keepout definition, which are implemented as functions within the module. The inputs and outputs for these functions are represented schematically in Fig. 2.

The observatory orbit plays a key role in determining which of the target stars may be observed for planet finding at a specific time during the mission lifetime. The observatory definition module's orbit function takes the current mission time as input and outputs the observatory's position vector. The position vector is standardized throughout the modules to be referenced to a heliocentric equatorial frame at the J2000 epoch. The observatory's position vector is used in the keepout definition task and target list module to determine which of the stars from the star catalog may be targeted for observation at the current mission time.

The duty cycle determines when during the mission timeline the observatory is allowed to perform planet-finding operations. The duty cycle function takes the current mission time as input and outputs the next available time when exoplanet observations may begin or resume, along with the duration of the observational period. The outputs of this task are used in the survey simulation module to determine when and how long exoplanet finding and characterization observations occur. The specific implementation of the duty cycle function can have significant effects on the science yield of the mission. For example, if the observing program is predetermined, such that exoplanet observations can only occur at specific times and last for specific durations, this significantly limits the observatory's ability to respond dynamically to simulated events, such as the discovery of an exoplanet candidate. This can potentially represent a suboptimal utilization of mission time, as it may prove to be more efficient to immediately spectrally characterize good planetary candidates rather than attempting to reobserve them at a later epoch. It also limits the degree to which follow-up observations can be scheduled to match the predicted orbit of the planet. Alternatively, the duty cycle function can be implemented to give the exoplanet observations the highest priority, such that all observations can be scheduled to attempt to maximize dynamic completeness[2] or some other metric of interest.

The keepout definition determines which target stars are observable at a specific time during the mission simulation and which are unobservable due to bright objects within the field of view, such as the sun, moon, and solar system planets. The keepout volume is determined by the specific design of the observatory and, in certain cases, by the starlight suppression system. For example, in the case of external occulters, the sun cannot be within the 180 deg annulus immediately behind the telescope (with respect to the line of sight) as it would be reflected by the starshade into the telescope. The keepout definition function takes the current mission time and star catalog module output as inputs and outputs a list of the target stars that are observable at the current time. It constructs position vectors of the target stars and bright objects, which may interfere with observations with respect to the observatory. These position

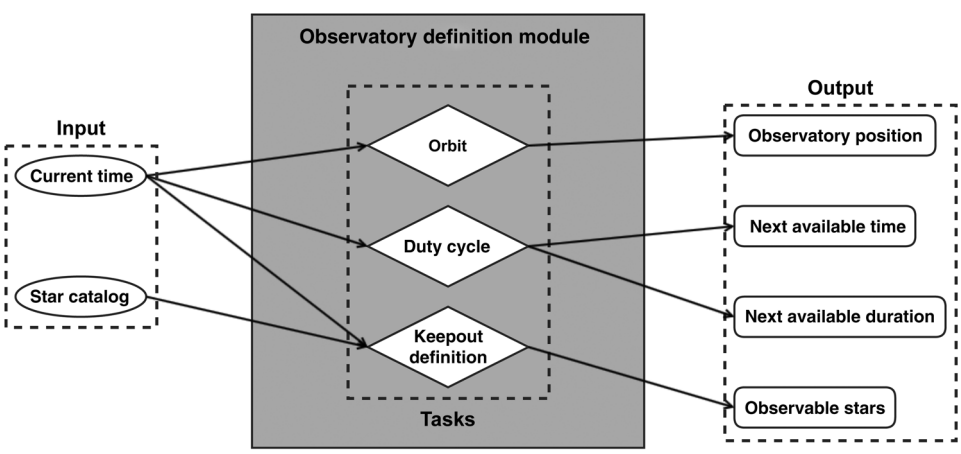

Fig. 2 Depiction of observatory definition module including inputs, tasks, and outputs.

vectors are used to determine if bright objects are in the field of view for each of the potential stars under exoplanet finding observation. If there are no bright objects obstructing the view of the target star, it becomes a candidate for observation in the survey simulation module.

The observatory definition also includes the target transition time, which encodes the amount of overhead associated with transitioning to a new target before the next observation can begin. For missions with external occulters, this time includes both the transit time between targets as well as the time required to perform the fine alignment at the end of the transit. For internal coronagraphs, this includes the settling time of the telescope to reach the bus stability levels required by the active wavefront control system. These may all be functions of the orbital position of the telescope and may be implemented to take into account thermal effects when considering observatories on geocentric orbits. This overhead calculation does not include any additional time required to reach the instrument's contrast floor, which may be a function of target brightness, and is encoded separately in the optical system description.

In addition to these functions, the observatory definition can also encode finite resources that are used by the observatory throughout the mission. The most important of these is the fuel used for stationkeeping and repointing, especially in the case of occulters that must move significant distances between observations. We could also consider the use of other volatiles such as cryogens for cooled instruments, which tend to deplete solely as a function of mission time. This module also allows for detailed investigations of the effects of orbital design on the science yield, e.g., comparing the baseline geosynchronous 28.5 deg inclined orbit for WFIRST-AFTA (Ref. 25) with an alternative L2 halo orbit also proposed for other exoplanet imaging mission concepts.[26]

2.1.5 *Planet physical model*

The planet physical model module contains models of the light emitted or reflected by planets in the wavelength bands under investigation by the current mission simulation. It uses physical quantities sampled from the distributions defined in the planet population, including planetary mass, radius, and albedo, along with the physical parameters of the host star stored in the target list module, to generate synthetic spectra or band photometry, as appropriate. The planet physical model is explicitly defined separately from the population statistics to enable studies of specific planet types under varying assumptions of orbital or physical parameter distributions, i.e., evaluating the science yield related to Earth-like planets under different definitions of the habitable zone. The specific implementation of this module can vary greatly and can be based on any of the many available planetary albedo, spectra, and phase curve models.[27–33]

2.1.6 *Time*

The time module is responsible for keeping track of the current mission time. It encodes only the mission start time, the mission duration, and the current time within a simulation. All functions in all modules requiring knowledge of the current time call functions or access parameters implemented within the time module. Internal encoding of time is implemented as the time from mission start (measured in days). The time module also provides functionality for converting between this time measure and standard measures, such as Julian Day Number and UTC time.

2.1.7 *Rules*

The rules module contains additional constraints placed on the mission design not contained in other modules. These constraints are passed into the survey simulation module to control the simulation. For example, a constraint in the rules module could include prioritization of revisits to stars with detected exoplanets for characterization when possible. This rule would force the survey simulation module to simulate observations for target stars with detected exoplanets when the observatory module determines those stars are observable.

The rules module also encodes the calculation of integration time for an observation. This can be based on achieving a predetermined signal-to-noise (SNR) metric (with various possible definitions) or via a probabilistic description as in Ref. 34. This requires also defining a model for the background contribution due to all astronomical sources and especially due to zodiacal and exozodiacal light.[5]

The integration time calculation can have significant effects on science yield—integrating to the same SNR on every target may represent a suboptimal use of mission time, as could integrating to achieve the minimum possible contrast on very dim targets. Changing the implementation of the rules module allows direct exploration of these tradeoffs.

2.1.8 *Postprocessing*

The postprocessing module encodes the effects of postprocessing on the data gathered in a simulated observation and the effects on the final contrast of the simulation. In the simplest implementation, the postprocessing module does nothing and simply assumes that the attained contrast is some constant value below the instrument's designed contrast—that postprocessing has the effect of uniformly removing background noise by a predetermined factor. A more complete implementation actually models the specific effects of a selected postprocessing technique, such as locally optimized combination of images (LOCI)[35] or Karhunen–Loève Image Projection (KLIP),[36] on both the background and planet signal via either processing of simulated images consistent with an observation's parameters or by some statistical description.

The postprocessing module is also responsible for determining whether a planet detection has occurred for a given observation, returning one of four possible states—true positive (real detection), false positive (false alarm), true negative (no detection when no planet is present), and false negative (missed detection). These can be generated based solely on statistical modeling as in Ref. 34, or can again be generated by actually processing simulated images.

2.2 *Simulation Modules*

The simulation modules include target list, simulated universe, survey simulation, and survey ensemble. These modules perform tasks that require inputs from one or more input modules as well as calling function implementations in other simulation modules.

2.2.1 *Target list*

The target list module takes in information from the optical system description, star catalog, planet population description, and observatory definition input modules and generates the input target list for the simulated survey. This list can contain either all of the targets where a planet with specified parameter ranges could be observed[37] or a list of predetermined targets such as in the case of a mission that only seeks to observe stars where planets are known to exist from previous surveys. The final target list encodes all of the same information as is provided by the star catalog module.

2.2.2 *Simulated universe*

The simulated universe module takes as input the outputs of the target list simulation module to create a synthetic universe composed of only those systems in the target list. For each target, a planetary system is generated based on the statistics encoded in the planet population description module, so that the overall planet occurrence and multiplicity rates are consistent with the provided distribution functions. Physical parameters for each planet are similarly sampled from the input density functions. This universe is encoded as a list where each entry corresponds to one element of the target list and where the list entries are arrays of planet physical parameters. In cases of empty planetary systems, the corresponding list entry contains a null array.

The simulated universe module also takes as input the planetary physical model module instance, so that it can return the specific spectra due to every simulated planet at an arbitrary observation time throughout the mission simulation.

2.2.3 *Survey simulation*

The survey simulation module takes as input the output of the simulated universe simulation module and the time, rules, and postprocessing input modules. This is the module that performs a specific simulation based on all of the input parameters and models. This module returns the mission timeline—an ordered list of simulated observations of various targets on the target list along with their outcomes. The output also includes an encoding of the final state of the simulated universe (so that a subsequent simulation can start from where a previous simulation left off) and the final state of the observatory definition (so that postsimulation analysis can determine the percentage of volatiles expended and other engineering metrics).

2.2.4 *Survey ensemble*

The survey ensemble module's only task is to run multiple simulations. While the implementation of this module is not at all dependent on a particular mission design, it can vary to take advantage of available parallel-processing resources. As the generation of a survey ensemble is an embarrassingly parallel task—every survey simulation is fully independent and can be run as a completely separate process—significant gains in execution time can be achieved with parallelization. The baseline implementation of this module contains a simple looping function that executes the desired number of simulations sequentially, as well as a locally parallelized version based on IPython Parallel.[38]

3 Wide-Field Infrared Survey Telescope-Astrophysics Focused Telescope Assets Coronagraph Modeling

While the development of EXOSIMS is ongoing, we have already produced simulation results with the functionality out of which the baseline EXOSIMS implementation is being built. In this section, we present the results of some mission simulations for WFIRST-AFTA using optical models of coronagraph designs generated at the Jet Propulsion Laboratory (JPL) during the coronagraph downselect process in 2013, as well as post-downselect optical models of the hybrid lyot coronagraph (HLC)[39] generated in 2014.[40] It is important to emphasize that the instrument designs and mission yields shown here are not representative of the final coronagraphic instrument or its projected performance. All of the design specifics assumed in these simulations are still evolving in response to ongoing engineering modeling of the observatory as a whole and to best meet the mission science requirements.

These simulations are instead presented in order to highlight the flexibility of the EXOSIMS approach to mission modeling and to present two important use cases. In Sec. 3.1, we present mission yield comparisons for different instrument designs while all other variables (observatory, star catalog, planet models, etc.) are kept constant. The results from these simulations are most useful for direct comparisons between different instruments and to highlight particular strengths and weaknesses in specific designs. Ideally, they can be used to guide ongoing instrument development and improve the final design science yield. In Sec. 3.2, we investigate a single coronagraph design operating under varying assumptions on observatory stability and postprocessing capabilities. These simulations highlight how EXOSIMS can be used to evaluate a more mature instrument design to ensure good results under a variety of operating parameters. This section also demonstrates how to incorporate the effects of different assumptions in the presimulation optical system diffractive modeling.

In addition to HLC, the first set of optical models includes models for a shaped pupil coronagraph (SPC)[41] and a phase-induced amplitude apodization complex mask coronagraph (PIAA-CMC).[42] In the downselect process, the SPC and HLC were selected for further development with PIAA-CMC as backup. It should be noted that the HLC optical models in the first and second set of simulations shown here represent different iterations on the coronagraph design, and thus, different instruments.

The optical system description is implemented as a static PSF, throughput curve, and contrast curve based on the JPL optical models. Other values describing the detector, science instrument, and the rest of the optical train were chosen to match Ref. 43 as closely as possible. The integration times in the rules module are determined via modified equations based on Ref. 34 to achieve a specified false positive and negative rate, which are encoded as constants in the postprocessing module. Spectral characterization times are based on preselected SNR values (as in Ref. 1) and match the calculations in Ref. 43.

The star catalog is based on a curated database originally developed by Turnbull et al.,[4] with updates to stellar data, where available, taken from current values from the SIMBAD astronomical database.[12] Target selection is performed with a detection integration time cutoff of 30 days and a minimum completeness cutoff of 2.75%.[37] Revisits are permitted at the discretion of the automated scheduler,[3] and one full spectrum

is attempted for each target (spectra are not repeated if the full band is captured on the first attempt). The total integration time allotted is 1 year, spaced over 6 years of mission time with the coronagraph getting top priority on revisit observations.

3.1 *Comparison of Pre-downselect Coronagraph Designs*

As a demonstration of EXOSIMS's ability to compare different instrument designs for a single mission concept, we compare mission simulation results based on optical models of the pre-downselect SPC, HLC, and PIAA-CMC designs. As all of these represent preliminary designs that have since been significantly improved upon, and as our primary purpose here is to demonstrate the simulations' utility, we will refer to the three coronagraphs simply as C1, C2, and C3 (in no particular order). Table 1 lists some of the parameters of the three coronagraphs, including their IWAs and OWAs, their minimum and mean contrasts, and maximum and mean throughputs. Each design has significantly different operating characteristics in its region of high contrast (or dark hole). C3 provides the best overall minimum contrast and IWA, but has a more modest mean contrast, whereas C2 has the most stable and lowest mean contrast over its entire dark hole, at the expense of a larger IWA. C1 has the smallest angular extent for its dark hole, but maintains reasonably high throughput. C2 has a constant and very low throughput, while C3 has the highest throughput over its entire operating region. Finally, while C1 and C3 cover the full field of view with their dark holes, C2 only creates high-contrast regions in 1/3 of the field of view, thus requiring three integrations to cover the full field.

We consider five specific metrics for evaluating these coronagraph designs:

1. unique planet detections, defined as the total number of individual planets observed at least once;
2. all detections, defined as the total number of planet observations throughout the mission (including repeat observations of the same planets);
3. total visits, defined as the total number of observations;
4. unique targets, defined as the number of target stars observed throughout the mission;
5. full spectral characterizations, defined as the total number of spectral characterizations covering the entire 400 to 800 nm band. This does not include characterizations where the IWA or OWA prevents full coverage of the whole band. This number will always be smaller than the number of unique detections based on the mission rules used here.

While it is possible to use EXOSIMS results to calculate many other values, these metrics represent a very good indicator of overall mission performance. As it is impossible to jointly maximize all five—in particular, getting more full spectra or additional repeat detections is a direct trade-off to finding additional, new planets—these values together describe the pareto front of the mission phase space. At the same time, these metrics serve as proxies for other quantities of interest. For example, taken together, all detections and unique detections indicate a mission's ability to confirm its own detections during the course of the primary mission, as well as for possible orbit fitting to detected planets. The number of unique targets, compared with the input target list, determines whether a mission is operating in a "target-poor" or "execution time-poor" regime. The latter can be addressed simply by increasing the mission lifetime, whereas the former can only be changed with an instrument redesign. Finally, comparing the numbers of unique detections and full spectra indicates whether an instrument design has sufficient capabilities to fully characterize the planets that it can detect.

For each of the coronagraphs, we run 5000 full mission simulations, keeping all modules except for the optical description and postprocessing constant. In addition to the parameters and implementations listed above, our postprocessing module implementation assumes a static factor of either 10 or 30 in terms of contrast improvement due to postprocessing. That is, results marked 10× assume that the achieved contrast on an observation is a factor of 10 below the design contrast at the equivalent angular separation. All together, we generated 30,000 discrete mission simulations, in six ensembles. Mean values and 1σ standard deviations for our five metrics of interest for each ensemble are tabulated in Table 2, with the full probability density functions (PDFs) shown in Figs. 3–7.

From the tabulated values, we see that the three coronagraphs have fairly similar performances in terms of number of planets found and spectrally characterized. Overall, C2 is most successful at detecting planets, due primarily to the stability of its contrast over the full dark hole. However, because of the very low overall throughput, this does not translate into more spectral characterizations than the other two designs. C1 and C2 benefit more from the change from 10 to 30× contrast improvement due

Table 1 Parameters for coronagraphs studied in Sec. 3.1.

Name	Inner working angle[a]	Outer working angle[a]	Contrast		Throughput[b]		Sharpness[c]	Field of view portion[d]
			Min	Mean	Max	Mean		
C1	0.128	0.652	6.91e-09	1.60e-08	0.40	0.32	0.0142	1
C2	0.184	1.064	4.06e-09	7.06e-09	0.22	0.22	0.0138	1/3
C3	0.085	0.624	2.87e-09	2.94e-08	1.00	0.85	0.0143	1

[a]Inner and outer working angle in arcseconds at 550 nm.
[b]This is the throughput due to the coronagraph optics only.
[c]Sharpness is defined as $(\sum_i P_i^2)/(\sum_i P_i)^2$ for normalized point spread function P_i.
[d]The fraction of the field of view covered by the coronagraph's region of high contrast.

Table 2 Mean values and standard deviations of five performance metrics calculated from ensembles of mission simulations for the instruments described in Table 1.

		Unique detections[b]		All detections[c]		Full spectra[d]		All visits		Unique targets	
Name	Contrast factor[a]	μ	1σ	μ	1σ	μ	1σ	μ	1σ	μ	1σ
C1	10×	1.8	1.4	2.6	4.4	1.1	1.1	74.9	28.2	63.5	3.7
	30×	3.7	2.0	4.4	2.6	2.2	1.5	56.4	2.7	55.3	2.2
C2	10×	2.4	1.6	7.8	13.5	1.4	1.2	141.3	38.8	74.9	0.4
	30×	4.2	2.1	5.0	2.9	2.2	1.5	59.5	1.9	58.1	1.2
C3	10×	2.0	1.5	2.4	2.0	1.3	1.2	54.7	2.0	54.0	1.5
	30×	3.0	1.9	3.4	2.3	1.9	1.4	31.5	2.2	30.9	1.9

[a]Contrast improvement factor due to postprocessing.
[b]Number of individual planets detected one or more times.
[c]Total number of detections (including repeat detections of the same planets).
[d]Total number of planets where spectra can be obtained over the whole wavelength range (400 to 800 nm).

to postprocessing than does C3, which already has the deepest overall contrast, but whose contrast varies significantly over the dark hole. The largest differences in the metrics are the total number of observations. These illustrate the direct trade-off between acquiring spectra, which takes a very long time, and doing additional integrations on other targets. In cases such as C2 with only 10× contrast improvement, the spectral characterization times are typically so long that most targets do not stay out of the observatory's keepouts, so the mission scheduling logic chooses to do more observations rather than wasting time on impossible spectral integrations.

Turning to the figures of the full distributions for these metrics, we see that despite having similar mean values for unique planet detections, the full distributions of detections are quite different, leading to varying probabilities of zero detections. As this represents a major mission failure mode, it is very important to track this value, as it may outweigh the benefits of a given design. C1 with only 10× contrast gain does particularly poorly in this respect, with over 15% of cases resulting in no planets found. However, when a 30× gain is assumed, C1 and C2 end up having the lowest zero detection probabilities. We again see that the effects of even this simple postprocessing assumption are not uniform over all designs. This is due to the complicated interactions between each instrument's contrast curve and the assumed distributions of planetary parameters. In essence, if our priors were different (leading to different completeness values for our targets), then we would expect different relative gains for the same postprocessing assumptions. This is always a pitfall of these simulations and must always be kept in mind when analyzing the results. It should also be noted that there have been multiple iterations of all these coronagraph designs since

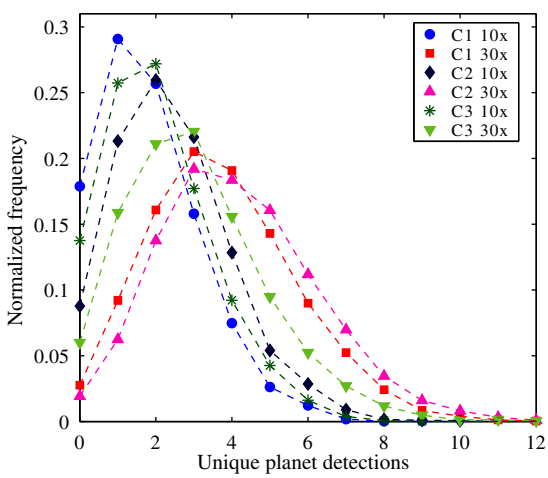

Fig. 3 Probability density functions (PDFs) of unique detections (number of individual planets, potentially with multiple planets about some targets, detected one or more times) for the coronagraph designs described in Table 1 assuming either a factor of 10 or 30 in postprocessing contrast gains. Of particular importance here is the probability of zero detections—all of the designs at 10× suppression and C1, in particular, have a significant (>5%) chance of never seeing a planet.

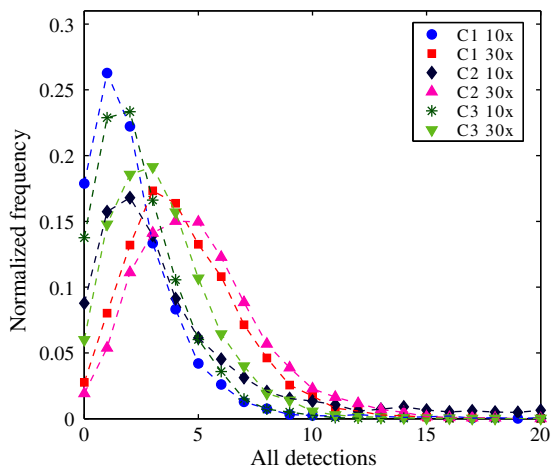

Fig. 4 PDF of all detections (including repeat detections) for instruments as in Fig. 3. Note that values of 15 or more typically represent a small number of easily detectable planets that are reobserved many times. Reobservations of a single target were capped at four successful detections in all simulations.

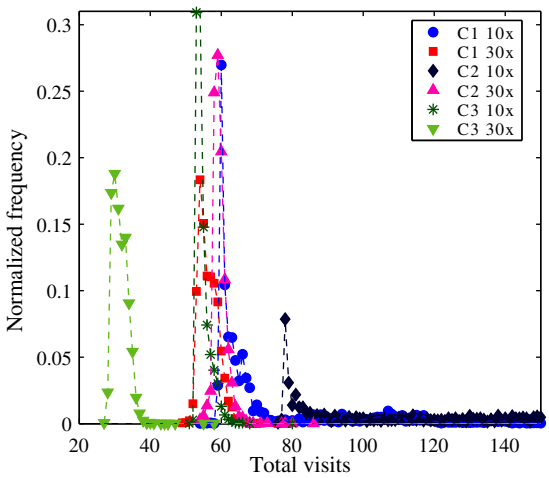

Fig. 5 PDF of total number of observations (including repeat observations of some targets) for instruments as in Fig. 3.

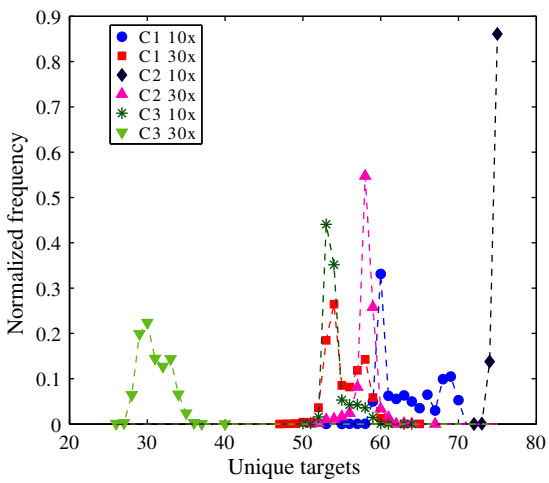

Fig. 6 PDF of unique targets observed for instruments as in Fig. 3. While all three instruments have fairly narrow distributions of this parameter, only C2 with 10× postprocessing gains is completely target limited.

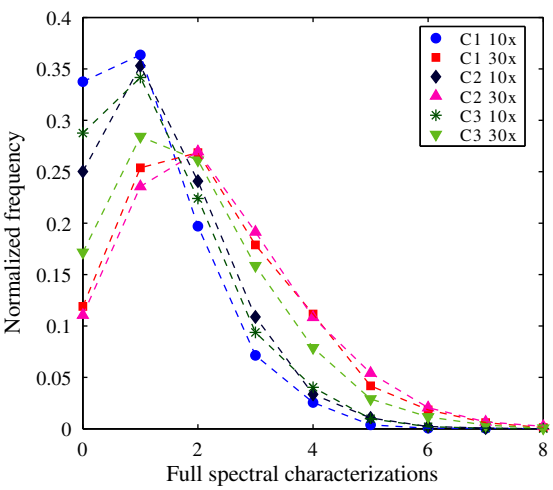

Fig. 7 PDF of number of spectra achieved over the whole band from 400 to 800 nm for instruments as in Fig. 3. C3 does comparatively well in this metric due to its lower inner working angle and high throughput.

downselect, resulting in significantly lower probabilities of zero detections, as seen in Sec. 3.2.

Another interesting feature is the very long right-hand tails of all detections and total visits distributions. These do not actually represent outliers in terms of highly successful missions, but rather typically imply the existence of one or a small number of very easy to detect planets. The logic of the scheduler allows the mission to keep returning to these targets for follow-up observations when it has failed to detect any other planets around the other targets in its list. This situation arises when the design of the instrument and assumptions on planet distributions leave a mission target limited. The distributions of unique targets show this limitation, with very narrow density functions for the actual number of targets observed for each instrument. In particular, Fig. 6 makes it clear that C2 with 10× postprocessing gains runs out of available targets. In order to combat this, the scheduler code prevents revisits to a given target after four successful detections of a planet around it. Finally, turning to Fig. 7, we see that all the three designs, regardless of postprocessing assumptions, have >10% probabilities of zero full spectral characterizations. C1 with 10× postprocessing gains fares most poorly with zero full spectra achieved in over one third of all cases.

Analysis of the survey ensembles also allows us to measure the biasing effects of the mission on the planet parameters of interest. As we know the input distributions of the simulation, we can think of these as priors and of the distribution of the observed planets as the posteriors. Figures 8 and 9 show the distributions of planetary mass and radius used in the simulations, respectively, along with the output distributions from the various coronagraph designs. The output distributions are calculated by taking the results of all of the simulations in each ensemble together, as the number of planets detected in each individual simulation is too small to produce an accurate distribution.

The input mass distribution shown here is derived from the Kepler radius distribution as reported in Ref. 44 and is calculated by assuming that this distribution is the same for all orbital periods and via an assumed density function.[7] The frequency spike seen at around 20 Earth masses is due to a poor overlap in the density functions used in this part of the phase space. This

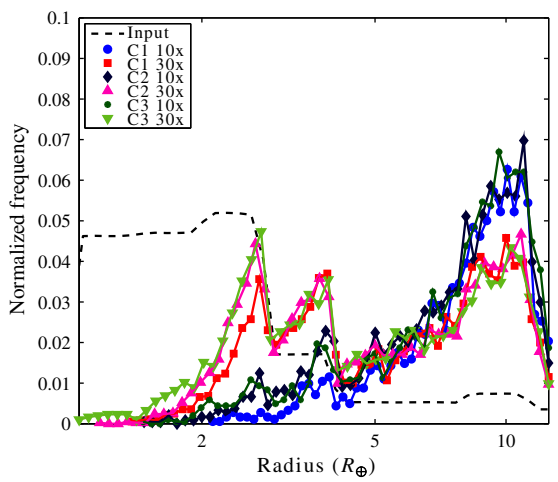

Fig. 8 Input and output distributions of planetary radius for instruments as in Fig. 3. The black dashed line represents the density function used in generating the planetary radii for the simulated planets in all simulations, while the other lines represent the distributions of planetary radii of the planets detected by each of the coronagraphs. The input distribution is based on the Kepler results reported in Ref. 44.

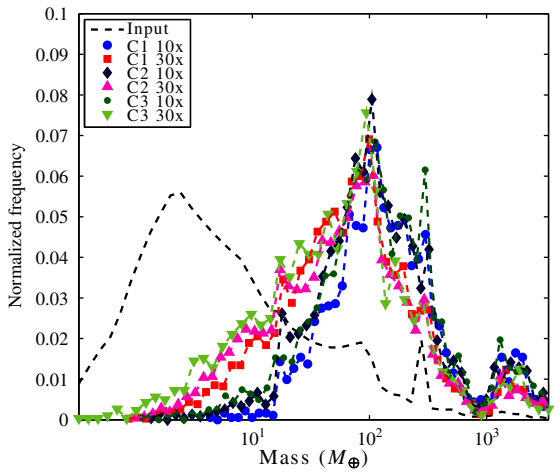

Fig. 9 Input and output distributions of planetary mass for instruments as in Fig. 3. The input mass distribution is derived from sampling the radius distribution shown in Fig. 8 and converting to mass via an assumed density function.

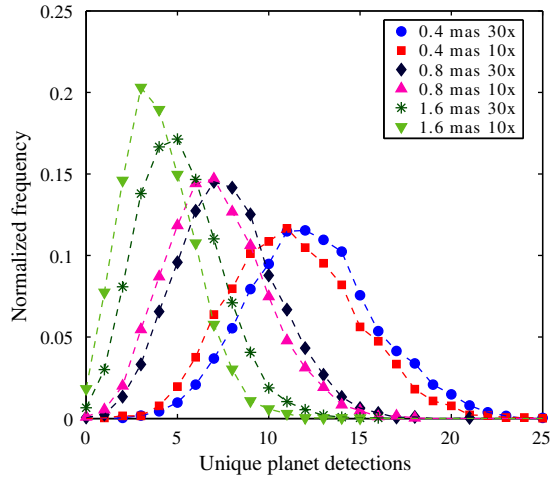

Fig. 10 PDF of unique planetary detections (number of individual planets, potentially with multiple planets about some targets, detected one or more times) for the post-downselect hybrid lyot coronagraph design, assuming either a factor of 10 or 30 in postprocessing contrast gains and telescope jitter of 0.4, 0.8, or 1.6 milliarcsec. It should be noted that the change in assumed postprocessing gain has a significantly smaller effect than the increased telescope jitter. We also note that the 1.6 milliarcsec jitter cases still have a small (0.6 to 1.8%) probability of never seeing a planet, whereas the 0.4 milliarcsec jitter ensembles do not contain a single simulation with zero plaents detected.

results in an equivalent spike in the posterior distributions, which slightly biases the results.

All of the instruments have fairly similar selection biases, although C1 and C3, which have smaller IWAs and higher throughputs, detect more low mass/radius planets. The effects of the instruments are readily apparent in all cases: lower radius planets, which are predicted to occur more frequently than larger radius ones, are detected at much lower rates.

3.2 *Comparison of Hybrid Lyot Coronagraph Parameters*

In this section, we present the results of survey ensemble analyses for a single instrument—a post-downselect HLC design—again assuming either 10 or 10× postprocessing gains, and assuming 0.4, 0.8, or 1.6 milliarcsec of telescope jitter. The jitter of the actual observatory will be a function of the final bus design and the operation of the reaction wheels, and its precise value is not yet known, which makes it important to evaluate how different levels of jitter may affect the achieved contrast and overall science yield. The jitter is built directly into the optical system model encoded in the optical system description module (see Krist et al., this volume, for details), while the postprocessing is treated as in Sec. 3.1.

As in Sec. 3.1, we run ensembles of 5000 simulations for each of the six cases considered, keeping all modules except for the optical description and postprocessing constant. The mean and 1σ of the five metrics of interest described in Sec. 3.1 are tabulated in Table 3, and the full PDFs for all metrics are shown in Figs. 10–14.

Table 3 Mean values and standard deviations of five performance metrics calculated from ensembles of mission simulations for the postdownselect hybrid lyot coronagraph with varying levels of assumed telescope jitter. Column definitions are as in Table 2.

		Unique detections[b]		All detections[c]		Full spectra[d]		All visits		Unique targets	
Jitter (milliarcsec)	Contrast factor[a]	μ	1σ	μ	1σ	μ	1σ	μ	1σ	μ	1σ
0.4	30×	12.4	3.5	14.0	4.4	9.5	3.2	47.6	4.3	45.2	4.0
	10×	11.4	3.5	12.5	4.2	6.2	2.6	31.9	2.4	30.4	1.9
0.8	30×	7.8	2.8	8.7	3.3	4.9	2.2	38.4	2.5	37.0	2.3
	10×	7.2	2.7	8.0	3.3	2.8	1.7	28.1	2.3	27.0	2.0
1.6	30×	5.1	2.3	5.7	2.7	1.4	1.2	31.6	1.6	30.8	1.4
	10×	4.0	2.0	4.4	2.4	1.9	1.4	44.9	2.2	44.1	2.2

[a]Contrast improvement factor due to postprocessing.
[b]Number of individual planets detected one or more times.
[c]Total number of detections (including repeat detections of the same planets).
[d]Total number of planets where spectra can be obtained over the whole wavelength range (400 to 800 nm).

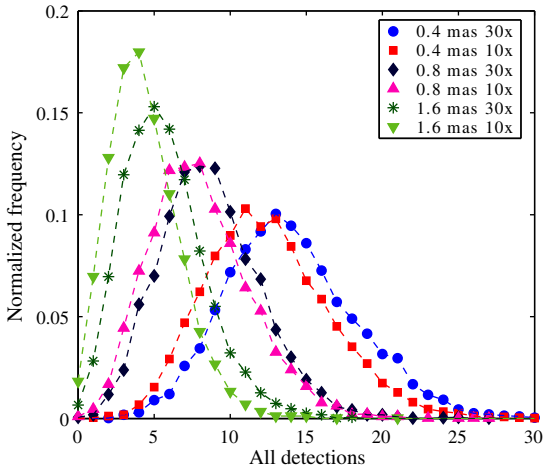

Fig. 11 PDF of total number of planetary detections (including repeat detections) for instruments as in Fig. 10. The trend here closely follows the one observed in the results for the unique detections metric.

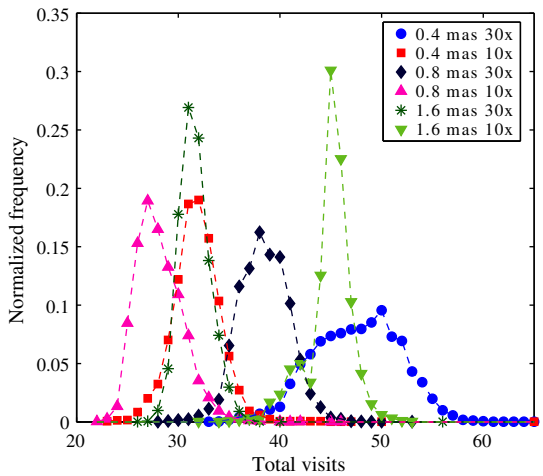

Fig. 12 PDF of total number of target observations (including repeat observations) for instruments as in Fig. 10. Here, the postprocessing improvement factor makes more of a difference than in the previous two figures, as more time must be devoted to spectral characterizations, limiting how much time is available for further observations.

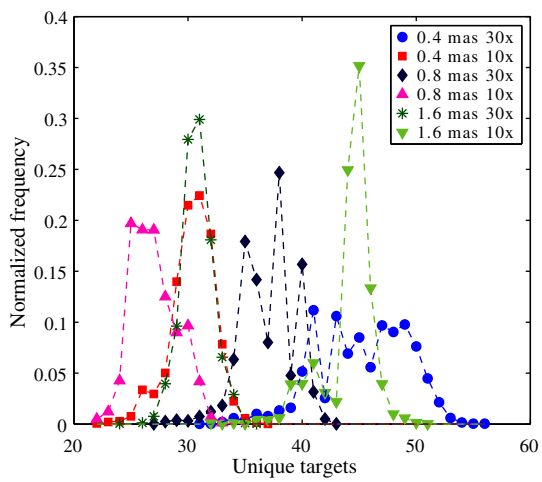

Fig. 13 PDF of unique targets observed for instruments as in Fig. 10. The trend here tracks closely to the one observed in the total visits metric and shows that this coronagraph design is not target limited in any of the studied cases.

One important observation made immediately obvious by these results is the relatively large effect of increased jitter versus the gains due to postprocessing. Tripling the assumed gain factor of postprocessing on the final achieved contrast has a significantly smaller effect on the number of detections, gaining only one unique detection, on average, as compared with halving the amount of telescope jitter, which increases the number of unique detections by >30%, on average. This shows us that the telescope jitter may be an effect that fundamentally cannot be corrected after the fact and, therefore, needs to be tightly controlled, with well-defined requirements set during mission design. Much of the current development effort for the project is focused on low-order wavefront sensing and control to mitigate these effects.[45,46]

We can also see significant improvements in the coronagraph design since the versions evaluated in Sec. 3.1, as the probability of zero planet detections is <2% in the case of the highest jitter level, and is well below 1% for all other cases. In fact, for both the 0.4 milliarcsec jitter ensembles, no simulations had zero detections, indicating a very low probability of complete mission failure for this coronagraph at these operating conditions.

Similar to the results of Sec. 3.1, the trend in the number of total visits does not simply follow those seen in the unique and total detection metrics, but is a function of both the number of detections and how much time is spent on spectral characterizations. We can see how the cases with the highest jitter and lowest postprocessing gains are pushed toward larger numbers of observations and unique targets, as they are able to achieve fewer full spectral characterizations, leaving them with additional mission time to search for new candidates. This is equally reflected in Fig. 14, where, despite the good performance seen in Fig. 10, all jitter levels have >5% chance of zero full spectra at the 10× postprocessing gain level, and only the 0.4 milliarcsec case at 30× gain has no instances of zero full spectra in its ensemble of results.

These metrics, taken together, clearly show that further optimization is possible via modification of mission rules, which were kept constant in all these ensembles. For example, the low numbers of spectral characterizations at higher jitter levels suggest that it may be worthwhile to attempt shallower

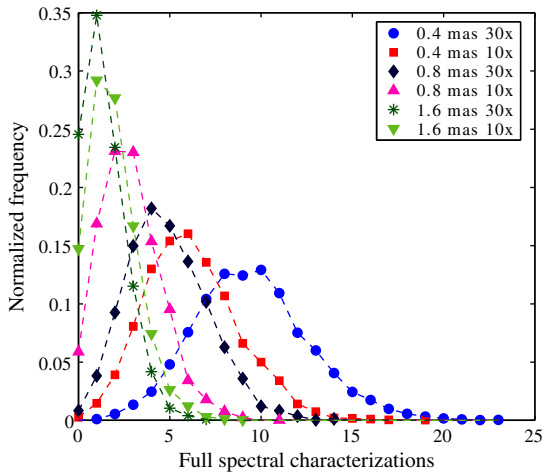

Fig. 14 PDF of number of spectra achieved over the whole band from 400 to 800 nm for instruments as in Fig. 10. In the worst case, there is a ~15% chance of not getting any spectra. Only the case of 0.4 milliarcsec jitter with 30× postprocessing gain has no simulations in its ensemble with zero full spectra achieved.

integrations in order to be able to make more total observations and potentially find a larger number of bright planets. This would bias the final survey results toward larger planets, but would increase the probability of spectrally characterizing at least some of the planets discovered. Alternatively, this may point to the desirability of investigating whether full spectral characterizations can be achieved for a small number of targets over the course of multiple independent observations.

4 Conclusions

We have presented the design details of EXOSIMS—a modular, open-source software framework for the simulation of exoplanet imaging missions with instrumentation on space observatories. We have also motivated the development and baseline implementation of the component parts of this software for the WFIRST-AFTA coronagraph, and presented initial results of mission simulations for various iterations of the WFIRST-AFTA coronagraph design.

These simulations allow us to compare completely different instruments in the form of early competing coronagraph designs for WFIRST-AFTA. The same tools also allow us to evaluate the effects of different operating assumptions, demonstrated here by comparing different assumed postprocessing capabilities and telescope stability values for a single coronagraph design.

As both the tools, the coronagraph, and mission design continue to mature, we expect the predictions presented here to evolve as well, but certain trends have emerged that we expect to persist. We have identified the portions of design space and telescope stability ranges that lead to significant probabilities of zero detections, and we expect instrument designs and observatory specifications to move away from these. We have also identified a mean number of new planetary detections for our particular assumed prior distributions of planetary parameters, that is consistent with the science definition team's mission goals for this instrument.

As we continue to both develop the software and improve our specific modeling of WFIRST-AFTA, we expect that these and future simulations will prove helpful in guiding the final form of the mission and will lay the groundwork for the analysis of future exoplanet imagers.

Acknowledgments

This material is based upon work supported by the National Aeronautics and Space Administration under Grant No. NNX14AD99G issued through the Goddard Space Flight Center. EXOSIMS is being developed at Cornell University with support by NASA Grant No. NNX15AJ67G. This research has made use of the SIMBAD database, operated at CDS, Strasbourg, France. The authors would like to thank Rhonda Morgan for many useful discussions and suggestions, as well as our reviewers, Wes Traub and Laurent Pueyo, who have significantly improved this work through their comments.

References

1. R. A. Brown, "Single-visit photometric and obscurational completeness," *Astrophys. J.* **624**, 1010–1024 (2005).
2. R. Brown and R. Soummer, "New completeness methods for estimating exoplanet discoveries by direct detection," *Astrophys. J.* **715**, 122 (2010).
3. D. Savransky, N. J. Kasdin, and E. Cady, "Analyzing the designs of planet-finding missions," *Pub. Astron. Soc. Pacific* **122**(890), 401–419 (2010).
4. M. C. Turnbull et al., "The search for habitable worlds: 1. The viability of a starshade mission," *Pub. Astron. Soc. Pacific* **124**(915), 418–447 (2012).
5. C. C. Stark et al., "Maximizing the exoearth candidate yield from a future direct imaging mission," *Astrophys. J.* **795**(2), 122 (2014).
6. D. Savransky and D. Garrett, "EXOSIMS," 2015, https://github.com/dsavransky/EXOSIMS
7. D. Savransky, "Space mission design for exoplanet imaging," *Proc. SPIE* **8864**, 886403 (2013).
8. K. L. Cahoy et al., "Wavefront control in space with MEMS deformable mirrors for exoplanet direct imaging," *J. Micro/Nanolith. MEMS MOEMS* **13**(1), 011105 (2014).
9. S. B. Shaklan et al., "Error budgeting and tolerancing of starshades for exoplanet detection," *Proc. SPIE* **7731**, 77312G (2010).
10. D. J. Denvir and E. Conroy, "Electron-multiplying CCD: the new ICCD," *Proc. SPIE* **4796**, 164–174 (2003).
11. A. Fowler and I. Gatley, "Demonstration of an algorithm for read-noise reduction in infrared arrays," *Astrophys. J.* **353**, L33 (1990).
12. M. Wenger et al., "The SIMBAD astronomical database—the CDS reference database for astronomical objects," *Astron. Astrophys. Suppl. Ser.* **143**(1), 9–22 (2000).
13. M. A. Perryman et al., "The hipparcos catalogue," *Astron. Astrophys.* **323**, L49–L52 (1997).
14. T. J. Henry, "The mass-luminosity relation from end to end," in *Proc. of the Workshop Spectroscopically and Spatially Resolving the Components of the Close Binary Stars*, pp. 159–165 (2004).
15. D. Savransky, E. Cady, and N. J. Kasdin, "Parameter distributions of Keplerian orbits," *Astrophys. J.* **728**(1), 66 (2011).
16. C. D. Dressing and D. Charbonneau, "The occurrence rate of small planets around small stars," *Astrophys. J.* **767**(1), 95 (2013).
17. J. Fortney, M. Marley, and J. Barnes, "Planetary radii across five orders of magnitude in mass and stellar insolation: application to transits," *Astrophys. J.* **659**(2), 1661 (2007).
18. A. Cumming et al., "The Keck planet search: detectability and the minimum mass and orbital period distribution of extrasolar planets," *Pub. Astron. Soc. Pacific* **120**(867), 531–554 (2008).
19. A. W. Howard et al., "The occurrence and mass distribution of close-in super-Earths, Neptunes, and Jupiters," *Science* **330**(6004), 653–655 (2010).
20. N. M. Batalha et al., "Planetary candidates observed by Kepler. III. Analysis of the first 16 months of data," *Astrophys. J. Suppl. Ser.* **204**(2), 24 (2013).
21. F. Fressin et al., "The false positive rate of Kepler and the occurrence of planets," *Astrophys. J.* **766**(2), 81 (2013).
22. E. A. Petigura, G. W. Marcy, and A. W. Howard, "A plateau in the planet population below twice the size of earth," *Astrophys. J.* **770**(1), 69 (2013).
23. J. McBride et al., "Experimental design for the Gemini Planet Imager," *Pub. Astron. Soc. Pacific* **123**(904), 692–708 (2011).
24. B. Macintosh et al., "First light of the Gemini Planet Imager," *Proc. Natl. Acad. Sci.* **111**(35), 12661–12666 (2014).
25. D. Spergel et al., "Wide-field infrared survey telescope-astrophysics focused telescope assets WFIRST-AFTA final report," *arXiv preprint arXiv:1305.5422* (2013).
26. D. Savransky et al., "Occulting ozone observatory science overview," *Proc. SPIE* **7731**, 77312H (2010).
27. J. B. Pollack et al., "Estimates of the bolometric albedos and radiation balance of Uranus and Neptune," *Icarus* **65**(2), 442–466 (1986).
28. M. S. Marley et al., "Reflected spectra and albedos of extrasolar giant planets. I. Clear and cloudy atmospheres," *Astrophys. J.* **513**(2), 879 (1999).
29. J. J. Fortney et al., "Synthetic spectra and colors of young giant planet atmospheres: effects of initial conditions and atmospheric metallicity," *Astrophys. J.* **683**(2), 1104 (2008).
30. K. L. Cahoy, M. S. Marley, and J. J. Fortney, "Exoplanet albedo spectra and colors as a function of planet phase, separation, and metallicity," *Astrophys. J.* **724**(1), 189 (2010).
31. D. S. Spiegel and A. Burrows, "Spectral and photometric diagnostics of giant planet formation scenarios," *Astrophys. J.* **745**(2), 174 (2012).
32. A. Burrows et al., "A nongray theory of extrasolar giant planets and brown dwarfs," *Astrophys. J.* **491**, 856 (1997).
33. A. Burrows, D. Sudarsky, and J. I. Lunine, "Beyond the t dwarfs: theoretical spectra, colors, and detectability of the coolest brown dwarfs," *Astrophys. J.* **596**, 587 (2003).

34. N. J. Kasdin and I. Braems, "Linear and Bayesian planet detection algorithms for the terrestrial planet finder," *Astrophys. J.* **646**, 1260–1274 (2006).
35. D. Lafrenière et al., "A new algorithm for point-spread function subtraction in high-contrast imaging: a demonstration with angular differential imaging," *Astrophys. J.* **660**(1), 770–780 (2007).
36. R. Soummer, L. Pueyo, and J. Larkin, "Detection and characterization of exoplanets and disks using projections on Karhunen-Loeve eigenimages," *Astrophys. J. Lett.* **755**(2), L28 (2012).
37. D. Savransky and N. J. Kasdin, "Design reference mission construction for planet finders," *Proc. SPIE* **7010**, 70101T (2008).
38. F. Perez and B. E. Granger, "IPython: a system for interactive scientific computing," *Comput. Sci. Eng.* **9**(3), 21–29 (2007).
39. J. Trauger et al., "Complex apodization lyot coronagraphy for the direct imaging of exoplanet systems: design, fabrication, and laboratory demonstration," *Proc. SPIE* **8442**, 84424Q (2012).
40. J. Krist, personal communication, JPL (2014).
41. N. Zimmerman et al., "A shaped pupil lyot coronagraph for WFIRST-AFTA," in *American Astronomical Society Meeting Abstracts*, 225 (2015).
42. E. Sidick et al., "Simulated contrast performance of phase induced amplitude apodization (PIAA) coronagraph testbed," *Proc. SPIE* **9143**, 91430W (2014).
43. W. A. Traub et al., "Science yield estimation for AFTA coronagraphs," *Proc. SPIE* **9143**, 91430N (2014).
44. F. Fressin et al., "The false positive rate of Kepler and the occurrence of planets," *Astrophys. J.* **766**(2), 81 (2013).
45. I. Poberezhskiy et al., "Technology development towards WFIRST-AFTA coronagraph," *Proc. SPIE* **9143**, 91430P (2014).
46. F. Shi, "Low order wavefront sensing and control for WFIRST-AFTA coronagraph," in *American Astronomical Society Meeting Abstracts*, 225 (2015).

Dmitry Savransky is an assistant professor in the Sibley School of Mechanical and Aerospace Engineering at Cornell University. He received his PhD from Princeton University in 2011, followed by a postdoctoral position at Lawrence Livermore National Laboratory, where he assisted in the integration and commissioning of the Gemini Planet Imager. His research interests include optimal control of optical systems, simulation of space missions, and image postprocessing techniques.

Daniel Garrett is a PhD student in the Sibley School of Mechanical and Aerospace Engineering at Cornell University. His research interests include dynamics and control theory, planetary science, and space exploration.