Uncertainty-aware Ground-based Telescope Observation Scheduling at ALMA

Justin Payan*[a], Noemi Barbagli[a], Ignacio Toledo[b], Rodrigo A. Carrasco[c], Sergio Martín[b,d], and Nihar B. Shah[a]

[a]Carnegie Mellon University, Pittsburgh, USA
[b]Joint ALMA Observatory, Santiago, Chile
[c]Pontificia Universidad Católica de Chile, Santiago, Chile
[d]European Southern Observatory, Garching bei München, Germany

ABSTRACT

Automated scheduling tools for ground-based observatories must balance immediate observing conditions with long-term scientific priorities and global scheduling constraints. Existing dynamic schedulers often fail to enforce such global constraints, and typically do not account for uncertainty in forecasts of environmental or operational availability, thereby exhibiting suboptimal performance. We address these limitations with **PULSAR: Predictive Uncertainty-aware Lookahead Scheduling with Adaptive Rebalancing**, an automated scheduling algorithm for dynamic telescope operations. PULSAR dynamically schedules observations, accounting for observatory priorities, global scheduling constraints, and uncertainty in forecasts. PULSAR introduces two key improvements: first, a dynamic weighting scheme that progressively increases penalties for global-constraint violations; and second, an *anticipatory* lookahead algorithm that plans over horizons of hours to days accounting for uncertainty in forecasts of weather and equipment availability. To our knowledge, PULSAR is the first dynamic telescope scheduling approach that includes these two factors.

We evaluate the effectiveness of PULSAR using simulations of the 2023-2024 observing cycle at the Atacama Large Millimeter/submillimeter Array (ALMA). Our evaluations reveal that while a simulation of ALMA's current dynamic scheduling algorithm violates constraints by $2.2 - 7.1\%$, PULSAR satisfies global constraints within a tolerance of just 0.1%. PULSAR also improves satisfaction of scientific goals by completing approximately 100 additional high-priority observations and 5 additional high-priority projects relative to ALMA's current scheduler. PULSAR is feasible for deployment, computing decisions in approximately 30 seconds on average. All code for PULSAR is available at https://github.com/justinpayan/PULSAR.

Keywords: Dynamic scheduling, Mathematical optimization, Stochastic optimization, ALMA, Simulation

1. INTRODUCTION

Dynamic scheduling algorithms must simultaneously meet three goals: optimize observatory objectives, meet short-term constraints on observability conditions, and meet long-term constraints by the end of the cycle. They select observations in real time, requiring decisions to be made in several minutes at most. The primary objective of an observatory is to complete observations that are important to the scientific community, as determined through peer review.[1] Observatories prefer to maximize the amount of time spent observing. They may also explicitly prioritize observations that are uniquely-enabled by newly-installed technology, such as setting high-frequency observations as high-priority after an upgrade that increases the range of observable frequencies. Observations must be performed in conditions that are good enough to obtain usable data, and the environmental requirements vary across observations. Some observatories also have constraints that apply at the end of an observing cycle, such as requiring specific time shares for participating partner organizations.[2–6]

Our algorithm, PULSAR, tackles two major problems with existing dynamic scheduling algorithms:

- Existing algorithms do not guarantee satisfaction of constraints which must hold by the end of a yearly cycle.
- Existing algorithms do not effectively incorporate uncertain forecasts of stochastic conditions such as weather and equipment availability, which can result in missed opportunities in achieving science goals.

*jpayan@andrew.cmu.edu

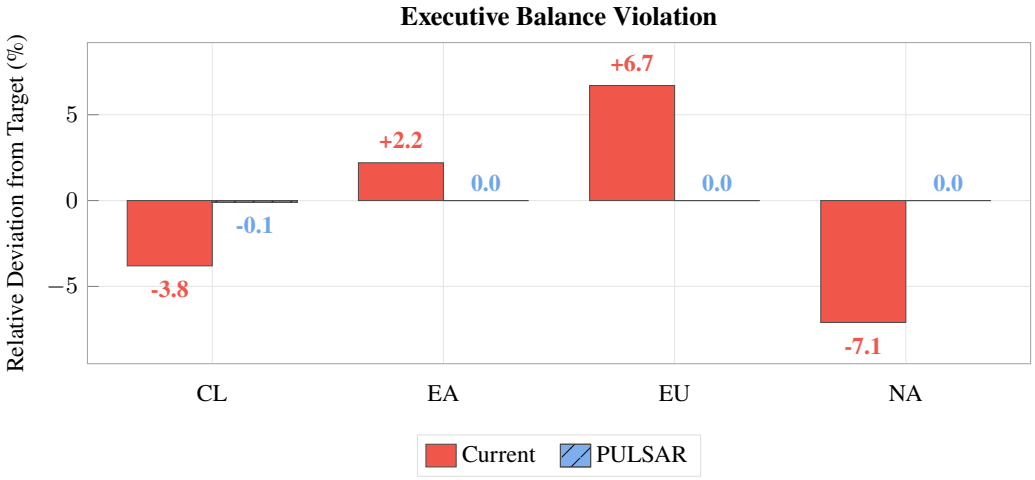

Figure 1. A comparison of executive balance deviations under a simulation of ALMA's current dispatch-based scheduler and under PULSAR.

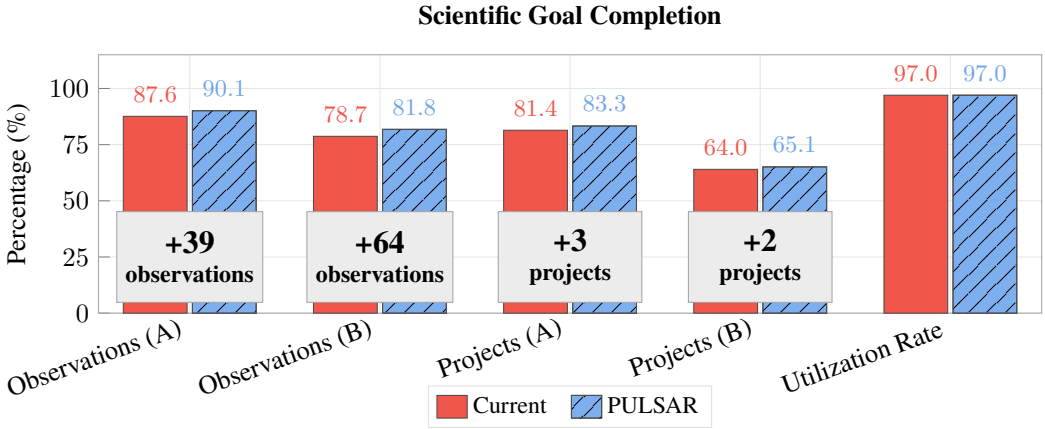

Figure 2. Comparison of observation completion rate, project completion rate, and utilization rate (100% minus percent idle time) under PULSAR as compared to a simulation of ALMA's current dispatch-based scheduler.

Satisfying long-term constraints. We are not aware of a published demonstration that an automated dynamic scheduling algorithm satisfies global time-sharing constraints at the end of an observing cycle. Time-share constraints apply at many telescopes that are built as partnerships, including ALMA,[2] Gemini,[3] the CFHT,[4] the HET,[5] and the upcoming CTAO.[6] These global constraints are *fundamental* to each telescope, since the time shares are based on funding contributions. If telescopes cannot reliably meet time allocation constraints, it would make it harder for partners to coordinate observations and maintain confidence in collaborative operations. However, public documentation typically describes these constraints as operational priorities or manually supervised queue-management goals, rather than reporting an automated dynamic scheduler that demonstrably satisfies them at the end of an observing cycle. Our PULSAR algorithm includes an adaptive rebalancing component that ensures we meet long-term/global constraints by the end of an observing cycle, without overemphasizing these constraints early in the cycle.

Consider the time-sharing constraint at ALMA. ALMA requires 33.75% of time for both North American and European projects, 22.5% of time for East Asian projects, and 10% of time for Chilean projects.[2] Their current dynamic scheduling approach fails to meet executive balance targets. Figure 1 plots the violations of executive balance constraints under a simulation of the $2023 - 2024$ observing cycle at ALMA. We find that ALMA's current automated dynamic scheduler violates executive balance constraints by between 2.2% and 7.1% relative deviation. This discrepancy is currently mitigated in practice with manual intervention. By contrast, PULSAR satisfies executive balance constraints within a tolerance of only 0.1%.

Improving efficiency by incorporating uncertain forecasts. Dynamic scheduling algorithms typically fall into two categories: *reactive* algorithms that apply heuristics derived from current conditions and do not create up-to-date, short-term plans,[7–12] and *rigid planning* algorithms that construct a plan over a planning horizon and then follow the plan with no modification or apply simple repair heuristics.[13–18] Reactive approaches fail to accurately anticipate the impact of immediate decisions on longer-term outcomes. For example, a reactive approach cannot foresee when an observation is about to go out of its observability window, or when bad weather is forecast in the next few hours. Reactive approaches can include heuristics reflecting longer-term concerns, but they do not have a principled method of comparing the relative values of immediate gains and long-term advantages. On the other hand, rigid planning approaches often consider a simplified model of the future, in which stochastic parameters (weather, equipment availability) are predicted with certainty. Rigid planning approaches fail when the true conditions deviate from the forecast conditions. Our PULSAR algorithm considers future conditions, but does not commit to a single plan. Instead, our algorithm explicitly plans for variation in future conditions, allowing us to make trade-offs dependent on the *relative likelihood* of future conditions instead of the single most likely outcome.

Our experiments — with results plotted in Figure 2 — find that compared to a simulation of ALMA's current dispatch-based scheduler, PULSAR completes approximately 100 additional high-priority observations and 5 additional high-priority projects. PULSAR exhibits an average decision time of approximately 30 seconds, making it feasible for deployment.

PULSAR is designed for any ground-based telescope accepting open proposals that are executed in autonomous mode or service mode. Such telescopes include the VLA, VLBA, and the upcoming NGVLA,[7] the GMRT,[13] CARMENES,[19] Gemini,[17,20] the CFHT,[8] the VLT,[20] Subaru,[21] the HET,[9,10] and the upcoming CTAO.[22,23] PULSAR is also applicable to ground-based telescopes that operate completely autonomously, such as the Joan Oró Telescope at the Montsec Astronomical Observatory (TJO-OAdM),[24] STELLA,[12,25,26] or the Vera Rubin Observatory.[27,28]

In the remainder of the paper, we describe the dynamic scheduling problem in more detail in Section 2, explain the PULSAR algorithm in Section 3, provide more details on our simulation setup in Section 4, present a dashboard for visualizing trade-offs between objective components and impacts of using different scheduling algorithms in Section 5, and review dynamic scheduling algorithms in Section 6.

2. DYNAMIC TELESCOPE SCHEDULING PROBLEM

At the beginning of an observing cycle, astronomers submit proposed projects. Through peer review, the observatory assigns each project an *importance weight*. Projects are also assigned *grades*; ALMA typically assigns approximately 100 projects an A grade and approximately 200 projects a B grade, while the rest receive a C grade, indicating they are filler projects. Projects are composed of *observations*, which are 30-minute to 2-hour tasks prepared for the telescope. One of the primary goals of an observing cycle is to complete high-priority projects (A-grade projects are the top priority, followed by B-grade projects) by completing all their component observations. This task is challenging because many observations have strict requirements on conditions in which they can execute.

The task of a dynamic scheduler is to select from available observations in real time, in such a way that scientific goals and constraints (such as executive balance) are met by the end of the cycle. Whenever the telescope becomes idle, the dynamic scheduler first ingests current conditions, including observation/project status, availability of equipment, any scheduled maintenance windows, weather data, and current forecasts for weather and/or equipment. Based on this information, the dynamic scheduler calculates the data quality expected from each observation in the queue, and ranks observations that have sufficient expected data quality. A purely automated telescope will then execute the top-ranked observation, or if there is an astronomer-in-the-loop they can incorporate the ranked list into their final observation decision. After the observation completes, the dynamic scheduler is invoked again. This loop continues until the end of the observing cycle, when we evaluate the performance of the observatory according to our objectives and constraints. The process is illustrated in Figure 3, which includes the relevant components for scheduling at ALMA in particular. For computational feasibility, we assume in this work that the cycle is decomposed into 30-minute increments.

Objectives. The observatory specifies objectives for measuring performance at the end of an observing cycle, and for guiding the decisions of the dynamic scheduler during the cycle. The observatory may also combine these objectives into a single, overall objective by adding them together. For the purposes of the current work, the objectives at ALMA are:

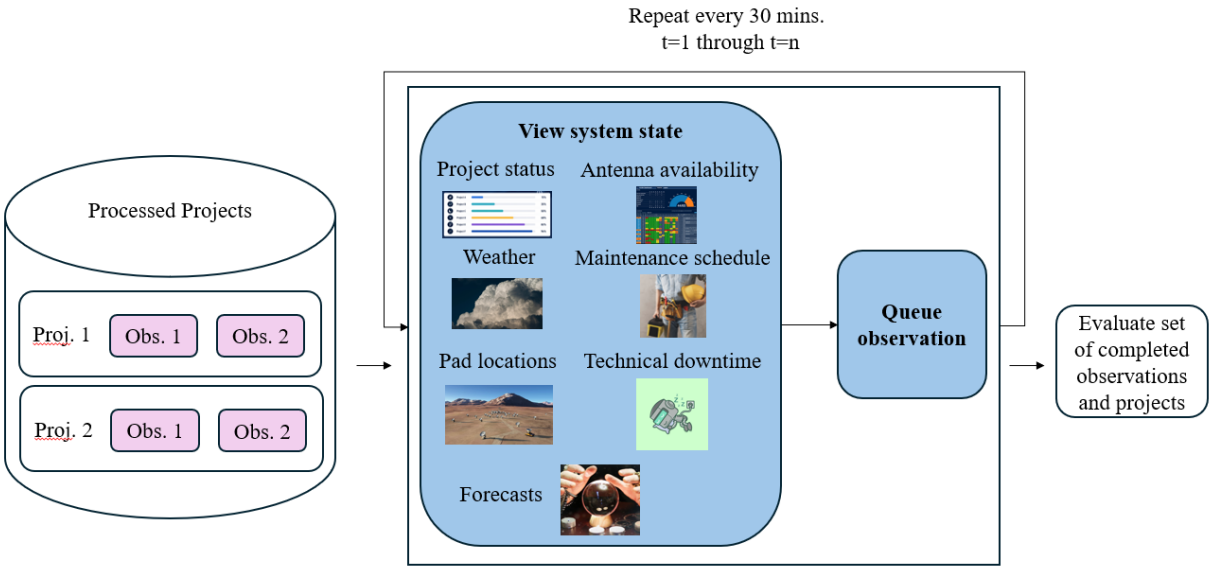

Figure 3. Illustration of the scheduling loop for ALMA.

- Observation completion: The sum of the importance weights of completed observations.
- Project completion: The sum of the importance weights of projects with all component observations completed.
- Idle time: The fraction of idle time is the total fraction of time usable for science observations which is not spent observing. We will track a value we call *utilization rate*, which is equal to $1 -$ (fraction of idle time).
- Executive balance violation: The *executive balance* constraint arises when multiple agencies/executives partner to construct and operate a telescope. The executives agree to an executive balance target, or an intended distribution of time allocated to each executive. ALMA has four executives: Europe (EU), North America (NA), East Asia (EA), and Chile (CL). EU and NA require 33.75% of time each, EA requires 22.5% of time, and CL requires 10% of time. Some deviation from the exact targets is inevitable, but ALMA aims to meet these targets as closely as possible each cycle. We measure the difference between the targets and the fraction of time actually observed under each executive.

We have specified four primary objectives – observation completion, project completion, utilization rate, and executive balance violation – though many others could also be included in our overall objective. The observatory decides how to weight each objective. Let the four-dimensional vector $\boldsymbol{\alpha}$ denote the weights assigned to each of the four objectives, and express the k^{th} objective listed above as $\textsc{obj}_k$. Let $\boldsymbol{\pi}$ denote a *schedule*, which is a realized assignment of observations to times throughout the cycle. We express our overall weighted objective as:

$$\textsc{obj}(\boldsymbol{\pi}) = \sum_{k=1}^{4} \alpha_k \textsc{obj}_k(\boldsymbol{\pi}). \tag{1}$$

Scheduling constraints. With the objective defined above, we can evaluate a completed schedule. However, we have not yet explained what constitutes a valid schedule. There are six constraints on schedules, which can be satisfied with certainty by the dynamic scheduler:

- Environmental conditions including weather and equipment availability are sufficient for an observation to reach a minimum level of data quality.
- Each observation can be scheduled at most once.
- Observations cannot overlap in time.

- Observations must end before the end of the cycle.
- Observations cannot overlap with scheduled maintenance.
- Observations cannot *start* during weather or technical downtime, but it is acceptable if an observation overlaps with weather or technical downtime after it is started.

The first constraint requires a bit more elaboration. We model the sequence of environment states as a matrix of random variables $\boldsymbol{x}$, where $\boldsymbol{x}_{:,t}$ denotes the full environment state at time t. At time t, the vector of random variables $\boldsymbol{x}_{:,t}$ is revealed, and the dynamic scheduler must make a scheduling decision based on the cycle history and the environment state $\boldsymbol{x}_{:,t}$. At ALMA the environment state consists of the weather (PWV and phase RMS), the availability of all antennas in each observing band, the locations of the antennas, and weather forecasts constructed by ALMA or provided by the National Oceanic and Atmospheric Administration (NOAA).[29]

We define a function $\mathcal{F}$ over schedules $\boldsymbol{\pi}$ and realizations of environmental states $\boldsymbol{x}$ such that for any schedule $\boldsymbol{\pi}$ and sequence of environmental conditions $\boldsymbol{x}$, $\mathcal{F}(\boldsymbol{\pi}, \boldsymbol{x}) = 1$ if and only if $\boldsymbol{\pi}$ satisfies all the constraints laid out in this section under $\boldsymbol{x}$. We can also define $\mathcal{F}$ over partial schedules and partial environmental state sequences, analogously.

Overall problem statement. Overall, we aim to select the optimal schedule $\boldsymbol{\pi}^*$ that maximizes the objective OBJ while respecting the constraints of $\mathcal{F}$ under the realization of the environment, $\boldsymbol{x}$. Formally, the problem is expressed as selecting

$$\boldsymbol{\pi}^* \in \arg\max_{\boldsymbol{\pi}} \textsc{obj}(\boldsymbol{\pi}) \quad \text{s.t. } \mathcal{F}(\boldsymbol{\pi}, \boldsymbol{x}) = 1. \tag{2}$$

If the environment states $\boldsymbol{x}_{:,t}$ are known with certainty for all $t \in [n]$, we can solve Equation (2) exactly using mathematical optimization libraries like Gurobi.[30] In real deployments, we will not know the environment states $\boldsymbol{x}$ in advance.

3. PULSAR: PREDICTIVE UNCERTAINTY-AWARE LOOKAHEAD SCHEDULING WITH ADAPTIVE REBALANCING

In this section, we introduce the PULSAR algorithm, which dynamically optimizes for Equation (2). PULSAR applies an adaptive penalty for long-term constraints such as executive balance, and leverages forecasts to plan 8 hours in advance. PULSAR takes into account the possibility that realized environmental conditions may deviate from forecasts. Although PULSAR is a heuristic method, it takes advantage of the fact that if environment conditions are known in advance, we can obtain the optimal solution to Equation (2) using optimization software such as Gurobi.[30]

Similarly to several dynamic scheduling algorithms in the literature,[14–16,31] PULSAR integrates both long-term and short-term planning. The long-term planning is important because simply planning for short-term outcomes ignores the long-term impacts of each scheduling decision. For example, if some observation has high availability for the entire cycle, we can de-prioritize it in early time steps.

However, existing short-term planning algorithms either do not plan ahead using forecasts at all, or plan using a single, fixed forecast. PULSAR samples multiple possible scenarios for the conditions over the next 8 hours, and creates plans for each scenario. It then ranks observations by the *average* performance achieved over the next 8 hours if each observation were scheduled in the present.

Satisfying long-term constraints. Schedulers do not need to meet long-term constraints like executive balance at every intermediate time step, only by the end of the observing cycle. When selecting observations in the short term, we apply a penalty for constraint violation that is low early in the cycle, but increases as the cycle continues.

For executive balance, we assume that low levels of violation may occur, but higher levels of violation become increasingly unacceptable. We therefore penalize according to the *squared* difference between the observed fraction of time for an executive and the target fraction. For any given (partial) schedule $\boldsymbol{\pi}$, this leads to a penalty of the form:

$$\begin{aligned}\mathsf{Penalty}_{\textsc{ExBal}}(\boldsymbol{\pi}) \doteq\ & (\%\text{ of time CL} - \text{CL target }\%)^2 \\ & + (\%\text{ of time EA} - \text{EA target }\%)^2 \\ & + (\%\text{ of time EU} - \text{EU target }\%)^2 \\ & + (\%\text{ of time NA} - \text{NA target }\%)^2\end{aligned}$$

We create a hyperparameter λ that controls how quickly the executive balance penalty increases. This parameter λ intuitively acts as an *executive balance looseness* parameter, that dictates how loosely the executive balance constraint is followed. If the observatory sets a higher value for λ, the executive balance penalty term becomes smaller at earlier time steps, making the final executive balance constraint hold more loosely. At every time step, we multiply the penalty term by (Fraction of cycle finished)$^\lambda$, and we multiply the other terms in the objective by $1 -$ (Fraction of cycle finished)$^\lambda$. With an appropriate setting of λ, we can ignore the long-term constraint penalty at the start of the cycle, but apply it much more heavily towards the end of the cycle.

Short-term lookahead with uncertain forecasts. PULSAR uses the framework of online stochastic combinatorial optimization (OSCO)[32] for short-term planning. As described in Section 2, each time we schedule an observation we make our decision using current conditions (observation/project completion, weather, equipment availability) as well as *forecasts* of future conditions.

We use forecasts that model the *distribution* of random environment variables in the future. The focus in this work is not building the most accurate forecasts, but rather to demonstrate improvements that can come from incorporating the full distribution over future random variables. For ALMA, we use the following forecasts:

- PWV: The National Centers for Environmental Prediction (part of the United States' National Oceanic and Atmospheric Administration or NOAA) publishes 5-day, point-estimate forecasts for PWV at many places on Earth.[29] They publish an updated forecast every 6 hours. At any time, we can access multiple pointwise forecasts for any future time in the next 8 hours, and construct a Gaussian distribution with mean and variance equal to the empirical mean and variance of the forecasts.
- Phase RMS: We build an unobserved components time series prediction model using historical phase RMS data.[33,34] This model considers a sliding window of recent phase RMS measurements, and returns means and standard deviations of Gaussian distributions for phase RMS at each future time, up to 8 hours in advance.
- Equipment availability: We assume that the availabilities and locations of all antennas remain unchanged for the next 8 hours.

The short-term lookahead follows the online stochastic combinatorial optimization approach.[32] At each decision point, PULSAR evaluates every observation that can be started immediately. For each candidate observation, PULSAR temporarily fixes that observation as the next scheduled observation, then samples several possible 8-hour futures for environmental conditions. For each sampled future, PULSAR solves a deterministic optimization problem over the lookahead horizon. The short-term plan for each scenario is generated using Gurobi. When Gurobi proves optimality within the allotted time, this gives the optimal feasible schedule for that sampled scenario; otherwise, PULSAR uses the best feasible schedule found within the time limit. Each resulting schedule is scored using the same objectives described earlier: observation completion, project completion, utilization rate, and executive balance violation penalty as described above. PULSAR then averages the scores obtained for each candidate observation across the sampled future scenarios. The observation with the highest average score is selected for execution.

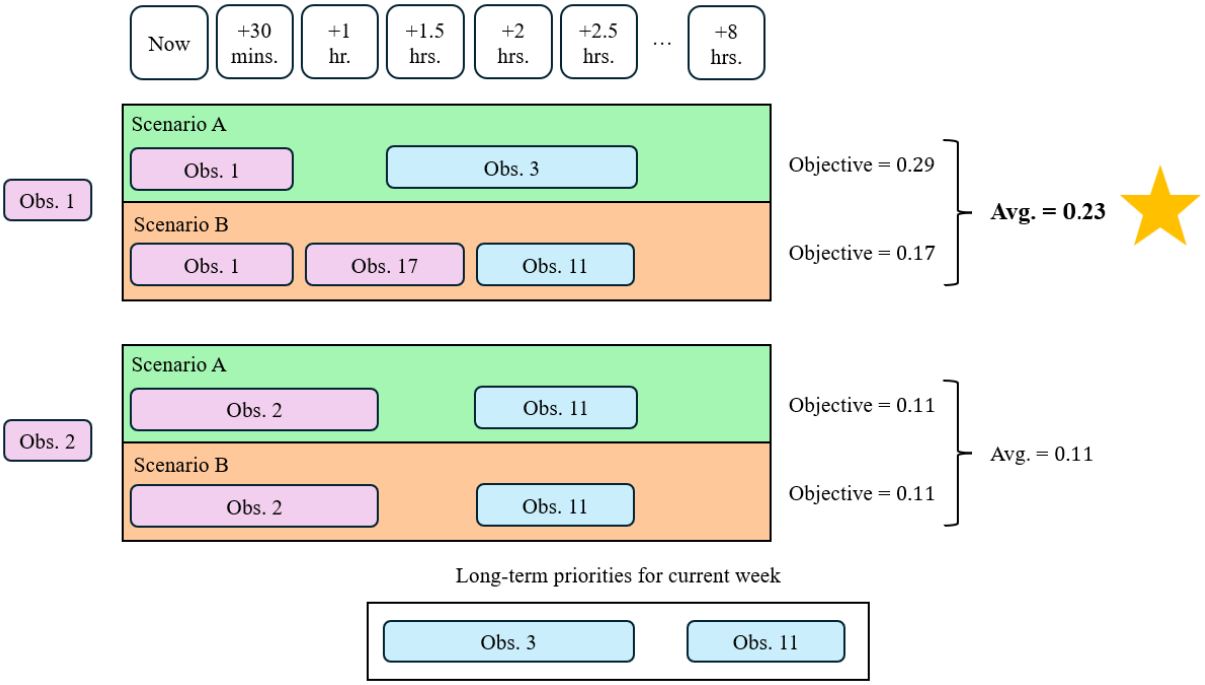

Figure 4. An example of PULSAR's short-term scheduling approach. We sample multiple scenarios of the next 8 hours, and optimize the schedule for each combination of currently schedulable observation and future scenario. Observations which have been prioritized by the long-term scheduler add bonus value to the short-term objective function. We queue the observation with the highest average short-term gain over all scenarios.

Long-term lookahead. In addition to the short-term lookahead, PULSAR also includes a long-term lookahead component. The long-term lookahead component solves a simplified version of Equation (2), which assigns observations to *weeks* rather than individual time slots. This separation of short-term and long-term planning is similar to other existing schedulers.[35] Our long-term scheduler is a modified version of a previously-developed long-term telescope scheduling model.[36]

One potential problem with following a long-term plan is that when conditions vary in real time, there may not be any available observations that are suggested by the long-term plan. In addition, if conditions are much better than expected, it may be possible to complete some important observations that were not originally planned. Thus rather than treating the long-term plan as fixed, we apply a *bonus* in the short-term planning objective for any observation which has been selected by the long-term schedule to be executed in current or previous weeks. This approach flexibly incentivizes adherence to the long-term plan, but allows the short-term scheduler to adapt to real-time conditions when necessary.

The long-term scheduler represents its recommendation as a set of observations. During the short-term lookahead, a candidate observation receives an additive bonus only if it is present in the recommended set. The bonus is grade-dependent: A-grade observations receive a stronger bonus than B-grade observations, while C-grade observations receive no long-term bonus. This mechanism encourages agreement with the long-term plan without making the plan a hard constraint.

Figure 4 shows an illustration of the process of selecting a single observation, with the bonus derived from long-term lookahead and the computation of outcomes for each currently-feasible observation. The feasible observation with the best expected outcome is selected for execution.

4. EXPERIMENTS

We execute ALMA's current dynamic scheduler and PULSAR on a simulation of ALMA's Cycle 10, which ran from October 2023 through September 2024. We present the resulting executive balance violations in Figure 1 and the other three objectives in Figure 2. This section provides more details about the experimental setup.

The simulation incorporates the projects, observations, maintenance windows, realized technical/weather downtimes, weather, antenna availabilities in each observing band, and NOAA PWV forecasts from ALMA's Cycle 10. We construct the

forecast model for phase RMS from historical values from $2017-2023$, excluding October-December of 2023, and then we issue forecasts for the Cycle 10 period using the forecast model. Consistent with the weights used in ALMA's current dynamic scheduler, we assign a weight of 1 to A-grade observations and projects, 0.4 to B-grade observations and projects, and -10.0 to C-grade observations and projects. We use long-term A-grade and B-grade bonuses of 10 and 1.5, respectively. We remove any observations which we know cannot be observed at any time during the cycle, as well as observations which have already been completed prior to the cycle. We then remove any projects which have no remaining observations to complete. Following this preprocessing, there are 1573 A-grade observations, 2078 B-grade observations, and 2465 C-grade observations. There are 156 A-grade projects, 172 B-grade projects, and 183 C-grade projects.

ALMA's current dispatch-based dynamic scheduling algorithm[31,37,38] ranks observations using a linear combination of scores computed from current conditions, without incorporating forecasts. It includes a score for how well the current conditions correspond to optimal conditions (conditions that are too good or too bad are both penalized), how close the current array comes to providing the desired angular resolution of the observation, a bonus proportional to the completion fraction of the observation's project and the number of times the observation must be executed[*], a bonus proportional to the overall rank of the observation's project, and a bonus promoting observations that are closer to transit. The algorithm does not take into account any prediction of future conditions or the overall set of observations and projects.

We execute PULSAR by rerunning the long-term scheduler every 2 weeks, and we use a short-term lookahead horizon of 8 hours. To ensure that executive balance constraints are met, we set α_4 much higher than the other α weights: $\alpha_1 = 0.002$, $\alpha_2 = 0.002$, $\alpha_3 = 0.001$, and $\alpha_4 = 0.995$. We apply a value of $\lambda = 10$ for the executive balance penalty weight schedule, as we found experimentally that this value results in meeting executive balance constraints. Setting $\lambda = 10$ causes the executive balance term to become relevant around 30% of the way through the cycle.

In our experiments, each OSCO decision samples five independent futures for the environmental conditions over the 8-hour lookahead horizon. For each candidate observation that can be started immediately and each sampled future, we solve a deterministic mixed-integer optimization problem over the lookahead horizon using Gurobi. We set Gurobi's relative optimality gap parameter to 0.0 and impose a 10-second time limit for each inner lookahead solve. Thus, whenever Gurobi proves optimality within the time limit, the exact optimal schedule for that sampled future is used; otherwise, PULSAR uses the best feasible solution found within the time limit. Experiments were run on a single node with 20 CPU cores and 160 GB of memory.

We find that ALMA's current dynamic scheduler selects observations in under 1 second, and PULSAR takes approximately 30 seconds to select an observation, on average. With this per-decision latency, a full year's simulation requires approximately $10-15$ minutes with ALMA's current scheduler, and approximately 1 day with PULSAR.

We report executive balance violations for the current dynamic scheduling algorithm compared to PULSAR in Figure 1. The executive balance targets are 0.3375 for EU, 0.3375 for NA, 0.225 for EA, and 0.1 for CL. We report the relative percent deviation from the target. For example, if we spend 9% of total time observing Chilean observations, the relative percent deviation from the target is 10%, since 0.01 is 10% of 0.1. Overall, we see that the current dynamic scheduler is unable to satisfy the executive balance constraint, while PULSAR satisfies the constraint within a tight tolerance.

We report observation completion rate, project completion rate, and utilization rate in Figure 2. The increase of 2.5% in A-grade observation completion accounts for 39 additional A-grade observations, and the 3.1% increase in B-grade completion accounts for 64 additional B-grade observations. The increase of 1.9% A-grade project completion and 1.1% B-grade project completion accounts for an additional 3 A-grade projects and 2 B-grade projects.

5. DASHBOARD

Because our overall objective is a weighted combination of multiple component objectives, we can trade off between the component objectives by changing the weight vector $\boldsymbol{\alpha}$. However, it is not always clear how the objectives interact. Some objectives are complementary, while others are conflicting. We develop a dashboard that enables astronomers and data scientists at ALMA to investigate trade-offs between the objectives of the observatory. The dashboard also demonstrates the impact of relaxing the executive balance constraint on scientific goals of observation/project completion and utilization rate. We show a screenshot of the dashboard in Figure 5.

[*]Although our model assumes observations can be executed at most once, when there are multiple identical observations we can group them together and track progress towards this set of observations.

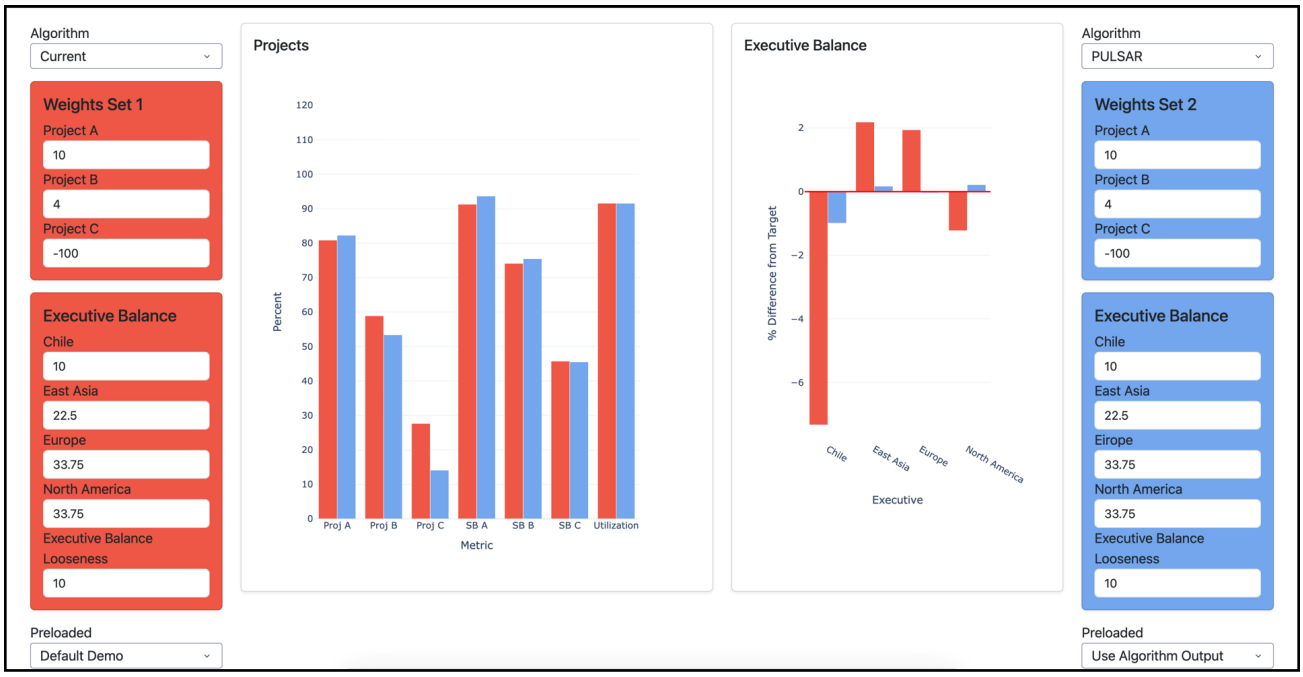

Figure 5. Screenshot of the dashboard comparing ALMA's current dynamic scheduler with PULSAR, using the same objective weights and executive balance looseness for each. Algorithms are color-coded, and it is straightforward to execute the current dynamic scheduler with different parameter values, or to upload new simulation results from any algorithm.

Users can either execute simulations directly in the dashboard, or upload a set of completed observations from a pre-computed simulation. The dashboard is currently configured to run simulations with Cycle 10 data, but can incorporate any historical data. The dashboard displays two sets of simulation outcomes at one time, with results color-coded for easy identification.

The dashboard allows setting project weights based on the grade assigned to the projects. The dashboard assumes all projects with the same grade have the same weight, and all observations within a project have the same weight as the overall project. Users can specify the fraction of time allocated to each executive under the executive balance constraint. Users may also specify the executive balance looseness λ that determines how quickly the executive balance penalty term increases in magnitude.

The dashboard currently visualizes A, B, and C grade project completion and observation completion (also called *scheduling block/SB* completion), utilization rate, and the percent violation of each executive time fraction.

6. RELATED WORK

Many observatories have implemented dynamic scheduling algorithms.[39] Most scheduling can be broken into long-term and short-term scheduling, once the project selection and grading has been completed (a challenging optimization problem in its own right, but considered out of scope of the current work).[35] Long-term scheduling assigns observations to longer time units such as days, weeks, or months, based on average historical conditions. Short-term scheduling queues observations in real time, taking into account current environmental and instrument conditions.

The primary contributions of the current work are to demonstrate a method for satisfying long-range constraints through long-term and short-term scheduling, and to incorporate uncertain forecasts into short-term scheduling. We are aware of only one approach that has introduced uncertainty over the future into short-term scheduling. Most other approaches either fail to plan ahead explicitly (a *reactive* approach), or plan assuming the forecasts are true to reality and re-plan if the weather deviates from the forecasts (a *rigid planning* approach).

One recent approach plans for uncertain future conditions.[40] The algorithm is developed for use at the VLT. However, this approach uses a simpler model than the model developed here, in which observations have only release dates and

deadlines, but otherwise can be executed at any time between the release dates and deadlines. The model assumes that some *nights* may not have conditions appropriate for observation, and the number of nights which will be observable is not known in advance. However, it is not accounting for finer-grained distinctions between observations with different requirements on stochastic conditions, nor accounting for variation from hour to hour.

We discuss some examples of reactive short-term scheduling. The system currently used by ALMA is one example of this approach.[31,38,41] Many other telescopes employ a similar dispatch-based, reactive short-term scheduler, including CARMENES,[19] the VLT,[11,42–46] STELLA,[12,26] and Gemini.[47] These approaches integrate long-term and short-term scheduling. Although these approaches consider many factors, none of them plans ahead more than one observation in advance.

The other category of short-term scheduler are those that plan ahead in the short-term, but assume a fixed plan over the planning horizon.[17,18,20,48–52] These approaches do not anticipate deviations between the forecasts and the realized conditions. A solution was proposed for ALMA that incorporates both long-term and short-term scheduling, but has not been deployed to date.[53] SPIKE[14,15,54,55] has been modified for use at several telescopes, including Subaru.[21,56] All of these short-term schedulers assume that stochastic conditions (e.g., weather and antenna availability) are known with certainty in advance, and the algorithms re-plan or repair the previous plan when conditions deviate from forecasts. Most of these approaches have not been explicitly evaluated against weather conditions deviating from forecasts, since the simulation environments in which they are tested also assume conditions are known with certainty in advance. Two recent papers focus on uncertainty in scheduling maintenance windows rather than weather.[50,51] As such, they do not *anticipate* changing conditions, though they do evaluate how their methods adapt to changing conditions after the fact. Our work explicitly anticipates deviations from forecasts, and we evaluate our algorithms against realistic conditions in which forecasts do not always match realized conditions.

7. CONCLUSION

We presented PULSAR, an algorithm that incorporates long-term and short-term planning using uncertain forecasts, to satisfy long-term constraints and maximize scientific productivity. We are currently running PULSAR alongside ALMA's existing dynamic scheduling model, with the aim of eventually integrating the model into daily operations. PULSAR and the dashboard we developed for evaluation are currently helping the observatory test hypotheses and make strategic decisions.

Our primary focus was not to create the best possible forecast models, but improving forecasting models would improve PULSAR's results. In particular, our model of antenna availability is quite naïve, and developing an improved model of antenna failures would improve the fidelity of the simulated scenarios used by the anticipatory lookahead. We hope this work encourages more observatories to consider variability in forecasts during short-term planning.

Acknowledgments

This work was supported by NSF CAREER 1942124. We are also grateful for a travel grant provided by the Joint ALMA Observatory (JAO) Visitor Program.

REFERENCES

[1] Carpenter, J. M., Corvillón, A., and Shah, N. B., "Enhancing peer review in astronomy: A machine learning and optimization approach to reviewer assignments for ALMA," *Publications of the Astronomical Society of the Pacific* **137**(3), 034501 (2025).

[2] "ISOpT 2026 ALMA users' policies, ALMA doc. 13.16 v1.0." https://almascience.nrao.edu/documents-and-tools/latest/alma-user-policies (2026). Edited by Takuma Izumi, Enrique Macias, Sergio Martín, and Catherine Vlahakis. Accessed: 2026-06-05.

[3] Gemini Observatory, "Policies regarding observing." https://www.gemini.edu/observing/policies (2025). Accessed: 2026-06-05.

[4] CFHT Scientific Advisory Council, "Report of the 71st meeting of the CFHT SAC." https://www.cfht.hawaii.edu/Science/SAC/SAC_website/reports/SAC_report_May07.html (2007). Accessed: 2026-06-05.

[5] Kelton, P. and Cornell, M., "Operations issues for the Hobby-Eberly Telescope," in [*Robotic Telescopes. Current Capabilities, Present Developments, and Future Prospects for Automated Astronomy*], **79**, 136 (1995).

[6] Cherenkov Telescope Array Observatory, "Groundbreaking ceremony marks the beginning of CTAO-South array construction in Chile." https://www.ctao.org/news/groundbreaking-ceremony-marks-the-beginning-of-ctao-south-array-construction-in-chile/ (2025). Accessed: 2026-06-05.

[7] National Radio Astronomy Observatory, "Scheduling." https://science.nrao.edu/observing/scheduling. Accessed: 2026-06-03.

[8] Canada-France-Hawaii Telescope, "Queued service observations." https://www.cfht.hawaii.edu/en/science/qso.php. Accessed: 2026-06-03.

[9] Shetrone, M., Cornell, M. E., Fowler, J. R., Gaffney, N., Laws, B., Mader, J., Mason, C., Odewahn, S., Roman, B., Rostopchin, S., Schneider, D. P., Umbarger, J., and Westfall, A., "Ten year review of queue scheduling of the Hobby-Eberly Telescope," *Publications of the Astronomical Society of the Pacific* **119**(855), 556–566 (2007).

[10] McDonald Observatory, "Hobby-Eberly telescope." https://mcdonaldobservatory.org/research/telescopes/HET. Accessed: 2026-06-03.

[11] Anderson, J. P., Sedaghati, E., Cikota, A., Behara, N., Bian, F., Otarola, A., and Mieske, S., "The optimisation of short-term scheduling of science observations at Paranal observatory (VLT and ELT)," in [*Observatory Operations: Strategies, Processes, and Systems X*], **13098**, 46–59, SPIE (2024).

[12] Granzer, T., "What makes an automated telescope robotic?," *Astronomische Nachrichten* **325**(6-8), 513–518 (2004).

[13] Gharote, M., Deshpande, A., Lodha, S., Kantharia, N., Wadadekar, Y., Katore, S., and Rao, A. P., "Automated telescope scheduling," in [*The Low-Frequency Radio Universe*], **407**, 438 (2009).

[14] Johnston, M. D., "Spike: AI scheduling for NASA's Hubble Space Telescope," in [*Sixth Conference on Artificial Intelligence for Applications*], 184–185, IEEE Computer Society (1990).

[15] Johnston, M. D. and Miller, G., "Spike: Intelligent scheduling of Hubble Space Telescope observations," *Intelligent Scheduling* **19**, 3–5 (1994).

[16] Garcia-Piquer, A., Morales, J. C., Ribas, I., Colomé, J., Guàrdia, J., Perger, M., Caballero, J. A., Cortés-Contreras, M., Jeffers, S. V., Reiners, A., Amado, P. J., Quirrenbach, A., and Seifert, W., "Efficient scheduling of astronomical observations - Application to the CARMENES radial-velocity survey," *Astronomy & Astrophysics* **604**, A87 (2017).

[17] Raaphorst, S., Miller, B., Nuñez, A., Troncoso, S., Rantakyro, F., Gimeno, G., and Walker, S., "Gemini automated scheduler software design." https://www.gemini.edu/files/software/Operations%20Development/Scheduler%20Design%20Review/5%20-%20FR_%20Gemini%20Automated%20Scheduler%20Software%20Design%20-%20Public.pdf (2022). Accessed: 2026-06-03.

[18] Baroch, D., Herrero, E., Gil, P., Domene, F., Rojas, A., and Ribó, N., "Autonomous scheduling at the TJO telescope," *Revista Mexicana de Astronomia y Astrofisica Serie de Conferencias (RMxAC)* **59**, 81–85 (2025).

[19] CARMENES Consortium, "CARMENES: Instrument." https://carmenes.caha.es/ext/instrument/index.html. Accessed: 2026-06-02.

[20] Amico, P., Campbell, R. D., and Christou, J. C., "Laser operations at the 8-10m class telescopes Gemini, Keck, and the VLT: lessons learned, old and new challenges," in [*Observatory Operations: Strategies, Processes, and Systems III*], **7737**, 77370A, SPIE (2010).

[21] Sasaki, T., Kosugi, G., Hawkins, R., Kawai, J. A., and Kusumoto, T., "Observation scheduling tools for Subaru Telescope," in [*Optimizing Scientific Return for Astronomy through Information Technologies*], **5493**, 367 – 372, SPIE (2004).

[22] Colomé, J., Colomer, P., Campreciós, J., Coiffard, T., De Oña, E., Pedaletti, G., Torres, D. F., and Garcia-Piquer, A., "Artificial intelligence for the CTA Observatory scheduler," in [*Observatory Operations: Strategies, Processes, and Systems V*], **9149**, 187–201, SPIE (2014).

[23] European Southern Observatory, "Cherenkov Telescope Array Observatory." https://www.eso.org/public/teles-instr/paranal-observatory/ctao/. Accessed: 2026-06-03.

[24] Colomé, J., Casteels, K., Ribas, I., and Francisco, X., "The TJO-OAdM Robotic Observatory: the scheduler," in [*Software and Cyberinfrastructure for Astronomy*], **7740**, 1176, SPIE (2010).

[25] Strassmeier, K. G., Granzer, T., Weber, M., Woche, M., Popow, E., Järvinen, A., Bartus, J., Bauer, S.-M., Dionies, F., Fechner, T., Bittner, W., and Paschke, J., "The STELLA Robotic Observatory on Tenerife," *Advances in Astronomy* **2010**(1), 970306 (2010).

[26] Granzer, T., Weber, M., and Strassmeier, K. G., "Three Years of Experience with the STELLA Robotic Observatory," *Advances in Astronomy* **2010**(1), 980182 (2010).

[27] Vera C. Rubin Observatory, "rubin_scheduler: Introduction." https://rubin-scheduler.lsst.io/introduction.html (2023). Accessed: 2026-06-03.

[28] Vera C. Rubin Observatory, "Scheduler procedures." https://obs-ops.lsst.io/v/DM-36406/Nighttime-Operations/Scheduler/index.html (2020). Accessed: 2026-06-05.

[29] NOAA National Centers for Environmental Information, "Global Forecast System (GFS)." https://www.ncei.noaa.gov/products/weather-climate-models/global-forecast (2020). Accessed: 2026-06-11.

[30] Gurobi Optimization, LLC, "Gurobi Optimizer Reference Manual," (2026).

[31] Avarias, J., Toledo, I., Espada, D., Hibbard, J., Nyman, L.-A., and Hiriart, R., "Towards a dynamical scheduler for ALMA: a science-software collaboration," in [*Software and Cyberinfrastructure for Astronomy IV*], **9913**, 1389–1402, SPIE (2016).

[32] Van Hentenryck, P. and Bent, R., [*Online stochastic combinatorial optimization*], The MIT Press (2006).

[33] Durbin, J. and Koopman, S. J., [*Time series analysis by state space methods*], Oxford University Press (UK) (2012).

[34] Fulton, C., Mendel, B., Quackenbush, P., Shedden, K., Sheppard, K., and Perktold, J., "statsmodels/statsmodels: Release 0.14.2," (2024).

[35] Colomé, J., Colomer, P., Guàrdia, J., Ribas, I., Campreciós, J., Coiffard, T., Gesa, L., Martínez, F., and Rodler, F., "Research on schedulers for astronomical observatories," in [*Observatory Operations: Strategies, Processes, and Systems IV*], **8448**, 469–486, SPIE (2012).

[36] Marcuello, C., *Desarrollo de una herramienta de apoyo para la programación estratégica de observaciones astronómicas del observatorio ALMA*, Master's Thesis, Pontificia Universidad Católica de Chile (2023).

[37] Privon, G., Popping, G., and Kawamura, A., "ALMA proposer's guide," (2024). Accessed October 7, 2024.

[38] Farris, A. R., "Scheduling simulation in ALMA," in [*Optimizing Scientific Return for Astronomy through Information Technologies*], **5493**, 302–312, SPIE (2004).

[39] Mora, M. and Solar, M., "A survey on the dynamic scheduling problem in astronomical observations," in [*Proceedings of the 3rd IFIP International Conference on Artificial Intelligence in Theory and Practice (IFIP AI)*], 111–120 (2010).

[40] Lacroix, T. R., Lemaire, P., Lagrange, A.-M., Milli, J., and Brauner, N., "Scheduling ground-based telescope observations with uncertain nights," (2026). Presented at the International Workshop on Planning & Scheduling for Space (IWPSS 2025).

[41] Clarke, D. and Avarias, J., "ALMA scheduling: It's dynamic!," *Astronomical Data Analysis Software and Systems XXI* **461**, 177 (2012).

[42] Patat, F. and Hussain, G. A., "Selecting and scheduling observing programmes at ESO," *Organizations, People and Strategies in Astronomy* **2**, 231–256 (2013).

[43] Alves, J., "Telescope time allocation tool," *The Messenger* **119**, 20–24 (2005).

[44] Rejkuba, M., Hainaut, O., Bierwirth, T., Pruemm, M., and Weiss, A., "Time allocation and long term scheduling of ESO telescopes at La Silla Paranal Observatory," in [*Observatory Operations: Strategies, Processes, and Systems X*], **13098**, 536–551, SPIE (2024).

[45] Silva, D., "Service mode scheduling: a primer for users," *The Messenger, No. 105, p. 18-24 (September 2001)* **105**, 18–24 (2001).

[46] Silva, D., "Service mode scheduling at the ESO Very Large Telescope," in [*Observatory Operations to Optimize Scientific Return III*], **4844**, 94–103, SPIE (2002).

[47] Puxley, P. J., "Execution of queue-scheduled observations with the Gemini 8-m telescopes," in [*Telescope Control Systems II*], **3112**, 234–245, SPIE (1997).

[48] Buffa, F., Serra, G., and Deiana, G. L., "Towards a dynamic scheduling system for the SRT," OAC Technical Report 67, INAF–Osservatorio Astronomico di Cagliari (2017).

[49] Catusse, N., Cambazard, H., Brauner, N., Lemaire, P., Penz, B., Lagrange, A.-M., and Rubini, P., "A branch-and-price algorithm for scheduling observations on a telescope.," in [*Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*], 3060–3066 (2016).

[50] Nakhjiri, N., Salamó, M., Sànchez-Marrè, M., and Morales, J. C., "A hybrid multi-start metaheuristic scheduler for astronomical observations," *Engineering Applications of Artificial Intelligence* **126**, 106856 (2023).

[51] Nakhjiri, N., Salamó, M., Sànchez-Marrè, M., Blum, C., and Carlos Morales, J., "Forgetful swarm optimization for astronomical observation scheduling," *IEEE Access* **12**, 171644–171661 (2024).

[52] Yáñez, J., "Optimization of telescope scheduling-algorithmic research and scientific policy," *Astronomy & Astrophysics* **403**(1), 357–367 (2003).

[53] Solar, M., Michelon, P., Avarias, J., and Garces, M., "A scheduling model for astronomy," *Astronomy and Computing* **15**, 90–104 (2016).

[54] Calvani, H., Civeit, T., England, M., Ake, T., Berman, A., Blair, W., Boyer, R., Caplinger, J., Kochte, M., Kruk, J., Roberts, B., and Suchkov, A., "The evolution of the FUSE mission planning system and operations," in [*SpaceOps 2008 Conference*], (2008).

[55] Muscettola, N. and Smith, S. F., "Investigations into generalization of constraint-based scheduling theories with applications to space telescope observation scheduling," NASA Contractor Report NASA-CR-202425, NASA Ames Research Center (1996).

[56] Sasaki, T., Kosugi, G., Kawai, J. A., Kusumoto, T., Koura, N., Hawkins, R., Kramer, L., Krueger, A. P., and Miller, G. E., "Observation scheduling scheme for the Subaru telescope," in [*Advanced Telescope and Instrumentation Control Software*], **4009**, 350–354, SPIE (2000).