

Astro2020 APC White Paper

Scheduling Discovery in the 2020s

Type of Activity: Ground Based Project Space Based Project
 Infrastructure Activity Technological Development Activity
 State of the Profession Consideration Other

Principal Author:

Name: Eric C. Bellm
Institution: University of Washington
Email: ecbellm@uw.edu
Phone: 206-685-2112

Co-authors:

Eric B. Ford (Penn State)
Aaron Tohuvavohu (Penn State)
Michael W. Coughlin (California Institute of Technology)

Endorsers:

Brett Morris (University of Bern)
Bryan Miller (Gemini Observatory)
Jennifer Sobeck (University of Washington)
Reed Riddle (California Institute of Technology)
Chuanfei Dong (Princeton University)
Peter Yoachim (University of Washington)

1 Overview

The 2020s will be the most data-rich decade of astronomy in history. On the ground and in space, powerful facilities like LSST, JWST, massive multi-object spectroscopic surveys, and a wide variety of smaller robotic and queue-based telescopes will be repeatedly scanning the sky. **However, without additional effort we are unlikely to realize the full scientific potential of our investments in these instruments.**

As the scale and complexity of our surveys and instruments increase, the problem of scheduling (which observations, in what order?) becomes more critical. First, prudence dictates maximizing the efficiency of facilities with high development costs and finite lifetimes. Second, key scientific projects—including identifying unseen populations of compact objects, understanding stellar binarity, and discovering and classifying rare classes of transients—require complex history-dependent observational sequences.

To date most scheduling of astronomical facilities has relied on very basic approaches: typically manual scheduling by human operators or simple “greedy” algorithms with basic objective functions. **We argue that these approaches are insufficient for the scientific needs of the 2020s.**

Thankfully, knowledge from fields such as Operations Research is beginning to percolate into astronomy. Surveys such as LCO, ALMA, ZTF, *Swift* and LSST are applying new algorithms to improve their efficiency and quantitative scientific throughput. However, much work remains to be done. **To maximize science in the 2020s, we must develop high-quality scheduling approaches, implement them as open-source software, and begin linking the typically separate stages of observation and data analysis.** The latter provides real-time feedback maximizing progress towards the scientific goal—the so-called “fifth paradigm”¹ of science (Szalay, 2019).

We provide an overview of key research directions as well as recommendations for scientists, facilities, and agencies to facilitate progress in the field.

2 Recommendations

2.1 Scientists

- Develop collaborations with scientists in operations research and related fields to help advance the state-of-the-art in astronomical scheduling and experimental design
- Develop formalisms connecting scientific results to the observational sequences required
- Incorporate fundamentals of experimental design into graduate curricula
- Continue to organize conference tracks at regular meetings like SPIE, ADASS, etc. as well as dedicated conferences² to share ideas and new approaches.

2.2 Surveys & Facilities

- Critically examine current scheduling and operations models and seek improvements
- Develop or adopt automated scheduling approaches
- Document and release scheduling tools as open-source packages to enable broad adoption
- Document scheduling decisions, so that future statistical analyses can mitigate biases³
- Adopt a shared schedule reporting schema to allow for efficient coordination and contextual scheduling decisions⁴

¹The first four paradigms are observation, analytic theory, computation and simulation, and data-intensive science (Hey et al., 2009).

²e.g., “Artificial Intelligence in Astronomy”, <https://www.eso.org/sci/meetings/2019/AIA2019.html>

³For example, which observations of a star were part of a baseline planet survey and which observations were done in order to expedite publication (regardless of whether decision is made by computer or human)?

⁴e.g., via the proposed IVOA Observation Locator Table Access Protocol, <http://www.ivoa.net/documents/ObsLocTAP/index.html>

2.3 Funding Agencies

- Formally evaluate plans for scheduling software and experimental design as part of mission/survey proposals
- Evaluate scheduling and operations models of current and forthcoming facilities
- Fund interdisciplinary efforts to develop new scheduling approaches and software, particularly that with applicability beyond single missions
- Fund contributions to general-purpose software libraries that schedulers and other high-level packages depend on.
- Recognize contributions to open-source software as a particularly high-value form of broader impact activity.

3 Impact on the Field

In the next decade, the scientific impact of a wide range of facilities can be enhanced by improvements in scheduling.

New large facilities with queue-scheduled Guest Observer programs run by time allocation committees will merit investments in efficient scheduling approaches appropriate to their staggering budgetary scale and finite lifetimes: these include JWST and the 30 m-class ground-based telescopes.

Large scale imaging surveys will continue to proliferate, with LSST, WFIRST, and Euclid only the most massive examples. Effective scheduling will be required to deliver cosmological constraints with minimal (and quantifiable) bias (e.g., Bechtol et al., 2019; Capak et al., 2019; Eifler et al., 2019; Kim et al., 2019), to detect and classify rare types of transients (e.g., Graham et al., 2019; Foley et al., 2019b), to inventory moving objects in our solar system (e.g., Milam et al., 2019; Chanover et al., 2019), and to trace stellar populations in our Galaxy and beyond (e.g., Kupfer et al., 2019; Rix et al., 2019; Price-Whelan et al., 2019). Smaller-aperture facilities may choose to adjust their own observing strategies to complement larger facilities or to explore focused niches for scientific returns too specialized for a general purpose survey.

Likewise, other non-electromagnetic facilities such as ground-based gravitational wave interferometers (e.g., LIGO) have observational selection effects (Chen et al., 2017) that, properly exploited via coordinated scheduling, can dramatically increase the serendipitous yield of detected multi-messenger sources (Tohuavohu et al., 2019).

Some time-domain surveys such as ZTF and LSST will be issuing real-time alert streams with millions of events nightly. Rapid, automated prioritization and robotic followup strategies first applied to gamma-ray burst followup can now be applied to a wide spectrum of time-domain events. Such followup will be required to fully understand rare and fast-evolving time domain events, such as kilonovae, “orphan” afterglows of gamma-ray bursts, rare types of stellar variables, and more.

Some of this followup will be undertaken by distributed networks of fully robotic telescopes (e.g.,

the Las Cumbres Observatory) which may be scheduled as a single entity (Lampoudi et al., 2015). However, an even greater opportunity and challenge is to realize the potential of heterogeneous telescope networks (e.g., Hessman, 2014) to stitch together disparate facilities in a uniform way. This can be useful for a few purposes. For example, using telescopes in both hemispheres can be important due to significant size of the sky localizations associated with short gamma-ray bursts and gravitational-wave sources. Often localizations have probability that is not fully accessible from a single location. It can also be useful to reimage the same localizations with multiple systems, both in different filters to measure an object’s color, as well at different times to differentiate between real transients and asteroids, as well as potentially measuring a change in luminosity. Coordinated scheduling between systems with significant differences between telescope setups, including their placement on the Earth, and their instrument configurations, including field of view, filters, typical exposure times, and limiting magnitudes, is desired.

Existing networks have begun to implement this network-level capability in open-source codebases (e.g., Coughlin et al., 2018)⁵. The Astronomical Event Observatory Network (AEON⁶) represents a current and more general step towards connecting multiple observatories through a common interface (Saunders et al., 2018, and see the Miller et al. 2019 APC white paper “Infrastructure and Strategies for MMA and Time Domain Follow-Up”), with a goal of enabling network-level scheduling of both imaging and spectroscopy for multiple teams on disparate ground-based facilities around the world.

Multi-wavelength and multi-team followup efforts, such as electromagnetic followup of gravitational wave triggers (e.g., Foley et al., 2019a; Sathyaprakash et al., 2019), present an additional layer of complexity. Coordinating disparate and perhaps competing PI teams using key shared facilities, each with their own TOO policies and scheduling constraints, is as much a political problem as an algorithmic one. But greater sharing of observations planned and undertaken could enable individual teams to pursue quantitatively rigorous approaches.

Massive multi-object spectroscopic surveys (DESI, PFS, SDSS-V, WEAVE, 4MOST, and more) will need to dynamically adapt their targeting strategy to account for changing observing conditions. The SDSS-V survey in particular has a strong time-domain component which will require effective scheduling (e.g., Rix et al., 2019; Kollmeier et al., 2019; Shen et al., 2019). With appropriate investment they could also opportunistically allocate spare fibers dynamically to follow up targets identified by time-domain imaging surveys.

Massive radial velocity surveys for exoplanets will demand rigorous scheduling approaches to answer key population questions (e.g., Ford et al., 2019; Bryson et al., 2019).

And finally, individual investigators with classically-scheduled nights will continue to seek to make the most effective use of the time they have available.

Across all of these diverse fields, improved scheduling methods can provide a range of benefits, including:

⁵<https://github.com/mcoughlin/gwemopt>

⁶<http://ast.noao.edu/data/aeon>

- Greater scientific productivity and reduced overheads
- Greater transparency and reproducibility
- Improved ability to rigorously simulate selection functions
- Potentially lower operations costs if replacing labor-intensive manual scheduling approaches (e.g., manually producing several versions of queue schedules for potentially different observing condition bands, most of which won't occur).
- The ability to dynamically respond to failures, changing observing conditions, and target of opportunity requests by repeatedly and regularly re-solving the scheduling problem (e.g., Lampoudi et al., 2015).

However, implementing state-of-the-art scheduling methods is currently not without its challenges:

- Initial development of new scheduling approaches requires specialized skills which are quite rare in astronomy.
- Changes to operations models and program prioritization may encounter user resistance unless the benefits (and any tradeoffs) are made clear.

The costs of implementing such scheduling approaches would be lowered if there were widely-available open-source scheduling software⁷ for the major classes of astronomical scheduling problems. The need to develop such algorithms and software forms the core of our recommendations.

4 Current Work

Current surveys are continuing to develop innovative scheduling approaches. These include both advances in the metrics or objective functions that surveys seek to maximize as well as new optimization algorithms.

4.1 Objective Functions

Many facilities on the ground and in space operate largely in a queue-based guest observer mode. Typically these observatories have disparate instruments and/or configurations, and the observing programs are selected on a periodic basis by a time allocation committee. Such facilities are well-served by optimization metrics that assess whether the proposals with higher TAC rank were executed in preference to those with lower rank, perhaps weighted by observing conditions. Lampoudi et al. (2015) and Solar et al. (2016) describe TAC-priority objective functions for the Las Cumbres Observatory robotic telescope network and ALMA, respectively.

Given their inherent multiplexing, large-scale imaging surveys often (but not exclusively) pursue multiple science goals simultaneously. They are accordingly well-served by objective functions that capture the rate at which they survey spatial volume (e.g., Bellm, 2016; Bellm et al., 2019) or

⁷See the E. Tollerud et al. APC whitepaper, “Sustaining Community-Driven Software for Astronomy in the 2020s,” for more on the scientific importance of open-source development and recommendations for sustaining it.

gain information (Tonry et al., 2018) or signal-to-noise.

Experiments that focus on one major science question can adopt metrics more specialized to their specific science goal. Bayesian objectives are particularly suitable in these cases, as they can maximize the expected change in information content of the posterior for the quantity being measured (Loredo, 2004; Ford, 2008).

4.2 Scheduling Algorithms

While the science metric to optimize will be specific to the survey or facility, in many cases the scheduling algorithm may be decoupled from the metric chosen. While astronomical scheduling problems have clear analogues in the optimization literature, most are different enough from standard textbook problems to prevent easy application of known solutions. For example, the problem of minimizing slew time when pointing a telescope can be cast as a Travelling Salesman Problem. However, for ground-based telescopes the earth’s rotation makes the time to slew between two points time-dependent. Moreover, the value of observing the fields at a specific time (the science metric) also is likely to change with time, for example as the airmass of the object changes (cf. Bellm et al., 2019).

Likewise, for space-based telescopes the slew-time between two points varies with time depending on many factors, including the exact dynamic Attitude Control System constraints of the spacecraft. Analogously, the value of observing fields at a specific time can change depending on the distance of the object from the Earth-limb at that time, for example (cf. Tohuvavohu, 2017).

Because many optimization problems have only heuristic solutions, the optimization literature is large, and many approaches have been explored by different astronomical projects (for a review, see Solar et al., 2016, and references therein). We highlight a handful of the wide range of approaches reported in the literature. In many cases hybrid approaches that apply different algorithms for short- and long-term scheduling are used (e.g., Colomé et al., 2010; Garcia-Piquer et al., 2014; Wetter et al., 2015).

“Greedy” or local search-based scheduling is quite common. The known shortcomings of greedy searches have led some authors to propose simple extensions that take into account astronomy-specific features like targets rising and setting (e.g., Denny, 2006; Rana et al., 2016). Others have adopted optimization approaches that heuristically attempt to identify better solutions than the current local optimum, including simulated annealing, genetic algorithms, Tabu Search, and Ant Colony Optimization (e.g., Colome et al., 2012; Garcia-Piquer et al., 2014).

In some cases generating a satisfactory observing sequence that avoids certain instrumental limitations is more important than maximizing an objective function. This suggests casting the scheduling problem in Constraint Satisfaction terms. A major example is the SPIKE system (Johnston & Miller, 1994), which has been used by the Hubble Space Telescope, Spitzer, Chandra, and the ground-based VLT. A related approach, TAKO (Timeline Assembler Keyword Oriented), has been utilized by *Swift* (2004–2018), *Fermi*, *Suzaku* and others. The *Swift* Scheduler (Tohuvavohu, 2017) and the open-source `astropplan` module (Morris et al., 2018) also employ

related approaches, though the *Swift* Scheduler combines this approach with a dynamic, fuzzy weighted priority scheme (Luo et al., 2003; Verfaillie & Jussien, 2005) and a flexible suite of optimization (objective) functions that yield a framework with generalities and automation capabilities more similar to Integer Linear Programming approaches.

Integer Linear Programming (ILP)-based approaches have been adopted by Las Cumbres Observatory (Lampoudi et al., 2015), ALMA (Solar et al., 2016), and ZTF (Bellm et al., 2019). These provide a more general framework for optimizing objective functions than constraint satisfaction while still enabling rigorous handling of constraints. Powerful commercial libraries are readily available⁸. Casting the scheduling problem in ILP terms enables rapid, regular reoptimization to update the schedule to respond to the success or failure of observations, changing weather conditions, etc.

Finally, some schedulers are based on artificial intelligence “agents” that are pre-trained to make decisions about observation sequences and then run on that trained model. Reinforcement learning approaches such as that of the LSST scheduler (e.g., Naghib et al., 2019) fall into this category. Reinforcement learning using deep neural networks has generated a great deal of popular press thanks to Google’s AlphaGo (Silver et al., 2016, 2017, 2018).

This profusion of solutions highlights the potential for breakthrough new approaches, but it also raises questions. How can we compare the performance of disparate systems applied to different scheduling problems? What are the performance and operational tradeoffs between different algorithmic approaches? Are there specific scheduling strategies that are provably superior for certain classes of problems found in astronomy?

5 Technology Drivers

In addition to the exciting new astronomical facilities which motivate us to seek improved scheduling (§3), in the 2020s we expect to see a continued technological improvements from outside astronomy that will impact the state of the art in scheduling.

In particular, rapid progress in Artificial Intelligence algorithms driven by commercial industry will continue, particularly in the near term around deep neural networks. Expensive optimization workflows will become more tractable due to a variety of factors. Access to scalable cloud computing will become easier and less expensive⁹. Cost per CPU core should continue to decrease, aiding optimization approaches which can be parallelized. Improved tooling for porting code to run on GPUs can yield even larger performance gains, and we are beginning to see application-specific integrated circuits such as Tensor Processing Units (TPUs) that are specifically developed for optimization workflows.

In conjunction with further algorithmic work (§6), these technological developments will bring currently intractable scheduling problems within reach within the next decade. For example, we might be able to ask, “Given all currently visible supernovae, what photometric data point should

⁸e.g., Gurobi, IBM CPLEX

⁹cf. the A. Smith et al. APC whitepaper “Astronomy Should be in the Clouds.”

we take next to most improve our constraint on the dark energy figure of merit?”

6 Strategic Plan

We prioritize five key areas for further development in the 2020s:

First, the development of a more general taxonomy of astronomical scheduling problems will illuminate mathematical commonalities and differences and guide algorithm selection.

Second, refinement of science-based metrics for optimization will enable surveys and facilities to better quantify their progress towards their scientific goals.

Third, development of more powerful scheduling algorithms will improve the overall throughput of existing facilities at modest cost.

Fourth, development of well-tested, well-documented open source implementations of scheduling algorithms will encourage reuse in the major problem domains, lowering the barrier to entry. The `astropy`-affiliated `astroplan` package provides one potential framework to which scheduling modules could be contributed.

Finally, connecting observations to data reduction and analysis will enable real-time optimization of a survey’s ultimate science goals. For example, an exoplanet RV survey might ask, “How does this observation improve my knowledge of the mass of a planet?” (Ford, 2008) or even better, “How does the observation improve my knowledge of the mass distribution of 2–4 R_{\oplus} planets with orbital periods of 30–300 d?” Research in the next decade will investigate how to make this “fifth paradigm” of science (Szalay, 2019) computationally feasible while preserving secondary science objectives and serendipitous discoveries.

References

- Bechtol, K., Drlica-Wagner, A., Abazajian, K. N., et al. 2019, in BAAS, Vol. 51, 207
- Bellm, E. C. 2016, PASP, 128, 084501
- Bellm, E. C., Kulkarni, S. R., Barlow, T., et al. 2019, PASP, 131, 068003
- Bryson, S., Bennett, D., Gaudi, S., et al. 2019, in BAAS, Vol. 51, 443
- Capak, P., Scolnic, D., & Davidzon, I. 2019, in BAAS, Vol. 51, 470
- Chanover, N., Wong, M. H., Greathouse, T., et al. 2019, in BAAS, Vol. 51, 340
- Chen, H.-Y., Essick, R., Vitale, S., Holz, D. E., & Katsavounidis, E. 2017, *The Astrophysical Journal*, 835, 31
- Colomé, J., Casteels, K., Ribas, I., & Francisco, X. 2010, in Proc. SPIE, Vol. 7740, Software and Cyberinfrastructure for Astronomy, 77403K
- Colome, J., Colomer, P., Guàrdia, J., et al. 2012, in Proc. SPIE, Vol. 8448, Observatory Operations: Strategies, Processes, and Systems IV, 84481L
- Coughlin, M. W., Tao, D., Chan, M. L., et al. 2018, *Monthly Notices of the Royal Astronomical Society*, 478, 692

- Denny, R. B. 2006, Dispatch Scheduling of Automated Telescopes,
<http://scheduler.dc3.com/Scheduler2006.pdf>, ,
- Eifler, T., Simet, M., Hirata, C., et al. 2019, in BAAS, Vol. 51, 418
- Foley, R., Alexander, K. D., Andreoni, I., et al. 2019a, in BAAS, Vol. 51, 295
- Foley, R., Bloom, J. S., Cenko, S. B., et al. 2019b, in BAAS, Vol. 51, 305
- Ford, E., Bedell, M., Ciardi, D. R., et al. 2019, Astro2020: Decadal Survey on Astronomy and Astrophysics, 2020, 466
- Ford, E. B. 2008, AJ, 135, 1008
- Garcia-Piquer, A., Guàrdia, J., Colomé, J., et al. 2014, in Proc. SPIE, Vol. 9152, Software and Cyberinfrastructure for Astronomy III, 915221
- Graham, M., Milisavljevic, D., Rest, A., et al. 2019, in BAAS, Vol. 51, 339
- Hessman, F. V. 2014, in The Third Hot-wiring the Transient Universe Workshop, ed. P. R. Wozniak, M. J. Graham, A. A. Mahabal, & R. Seaman, 95–101
- Hey, T., Tansley, S., & Tolle, K. 2009, The Fourth Paradigm: Data-Intensive Scientific Discovery (Microsoft Research)
- Johnston, M. D., & Miller, G. E. 1994, in Intelligent Scheduling, ed. M. Zweben & M. S. Fox (Morgan Kaufmann Publishers), 391
- Kim, A., Aldering, G., Antilogus, P., et al. 2019, in BAAS, Vol. 51, 140
- Kollmeier, J., Fuller, J., Gaensicke, B., et al. 2019, in BAAS, Vol. 51, 503
- Kupfer, T., Kilic, M., Maccarone, T., et al. 2019, in BAAS, Vol. 51, 188
- Lampoudi, S., Saunders, E., & Eastman, J. 2015, ArXiv e-prints, arXiv:1503.07170
- Loredo, T. J. 2004, in American Institute of Physics Conference Series, Vol. 707, Bayesian Inference and Maximum Entropy Methods in Science and Engineering, ed. G. J. Erickson & Y. Zhai, 330–346
- Luo, X., man Lee, J. H., fung Leung, H., & Jennings, N. R. 2003, Fuzzy Sets and Systems, 136, 151
- Milam, S., Hammel, H., Bauer, J. M., et al. 2019, in BAAS, Vol. 51, 327
- Morris, B. M., Tollerud, E., Sipőcz, B., et al. 2018, AJ, 155, 128
- Naghieb, E., Yoachim, P., Vanderbei, R. J., Connolly, A. J., & Jones, R. L. 2019, AJ, 157, 151
- Price-Whelan, A., Breivik, K., D’Orazio, D., et al. 2019, in BAAS, Vol. 51, 206
- Rana, J., Singhal, A., Gadre, B., Bhalerao, V., & Bose, S. 2016, ArXiv e-prints
- Rix, H.-W., Ting, Y.-S., Sippel, A., et al. 2019, in BAAS, Vol. 51, 104
- Sathyaprakash, B., Bailes, M., Kasliwal, M. M., et al. 2019, in BAAS, Vol. 51, 276
- Saunders, E. S., Bowman, M., Boroson, T., & Street, R. A. 2018, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 10704, Proc. SPIE, 107040Z
- Shen, Y., Anderson, S., Berger, E., et al. 2019, in BAAS, Vol. 51, 274
- Silver, D., Huang, A., Maddison, C. J., et al. 2016, Nature, 529, 484
- Silver, D., Schrittwieser, J., Simonyan, K., et al. 2017, Nature, 550, 354
- Silver, D., Hubert, T., Schrittwieser, J., et al. 2018, Science, 362, 1140
- Solar, M., Michelon, P., Avarias, J., & Garces, M. 2016, Astronomy and Computing, 15, 90
- Szalay, A. 2019, in American Astronomical Society Meeting Abstracts, Vol. 233, American Astronomical Society Meeting Abstracts #233, 400.01
- Tohuvavohu, A. 2017, arXiv e-prints, arXiv:1712.08111
- Tohuvavohu, A., Evans, P., Kennea, J., & Cenko, S. B. 2019, in American Astronomical Society Meeting Abstracts, Vol. 233, American Astronomical Society Meeting Abstracts #233, 210.08

- Tonry, J. L., Denneau, L., Heinze, A. N., et al. 2018, PASP, 130, 064505
- Verfaillie, G., & Jussien, N. 2005, Constraints, 10, 253
- Wetter, J., Akgun, Ö., Barker, A., et al. 2015, in 2015 IEEE 11th International Conference on e-Science, IEEE, 362–370

Astro2020 Science White Paper: Cosmology with a Space-Based Gravitational Wave Observatory .pdf. The low frequency, mHz GW spectrum is a new frontier for cosmology. We present the cosmological questions that can be addressed from observations in this regime. In this white paper we summarize the opportunity for discovery space in the coming decade of isolated brown dwarfs with planetary masses in young stellar associations within 150 pc. We suggest that next generation telescopes need to invest in characterizing young brown dwarfs in order to fully understand the atmospheres of sibling exoplanets. Over 200 data center white papers created by a dedicated research team, provide fact-based analysis to address and clarify today's challenges in planning, designing, operating, and maintaining data center facilities. Subjects cover power, cooling, DCIM, energy efficiency, operations & maintenance, new technologies, and more. Top Downloaded Papers. Learn about the subjects our customers are most interested in. These are the top downloaded white papers. Calculating Total Cooling Requirements for Data Centers. The Seven Types of Power Problems. Comparing UPS System Design Configurations. You will receive your welcome email soon describing the program and what to expect in the upcoming weeks. Enjoy! Your browser is out of date and has known security issues. ASTRO2020 APC White Paper. Orbiting Starshade: Observing Exoplanets at visible wavelengths with GMT, TMT, and ELT. Thematic Area: Planetary Systems. 8. Schedule: The Team X cost estimate was based on a launch in 2035. The assumed schedule was 94 months from Phase A through D. This allows for ~8 years of pre-phase A. 9. Cost estimates: The cost estimate was performed by Team X at JPL in May 2019. The base fiscal year was 2019, and the top-level cost was \$3.0 B in FY2019 dollars.