

USING FLEXIBLE EXECUTION, REPLANNING, AND MODEL PARAMETER UPDATES TO ADDRESS ENVIRONMENTAL UNCERTAINTY FOR A PLANETARY LANDER

Virtual Conference 19–23 October 2020

Daniel Wang¹, Joseph A. Russino¹, Connor Basich¹, Steve A. Chien¹

¹*Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive Pasadena, California 91109, USA, E-mail: firstname.lastname@jpl.nasa.gov*

ABSTRACT

When planning in unknown environments, a planner must ensure robustness to unknown phenomena and manage unpredictable variation in execution. Productivity in the face of these challenges requires an integrated approach to planning and execution that is capable of reacting to variation while still maximizing overall utility. We examine this problem in the context of the Europa Lander mission concept. We propose a planning and execution framework that responds to feedback using three techniques: (1) flexible execution, (2) periodic replanning, and (3) online model parameter and utility updates. The efficacy of each of these techniques is examined through simulation of the Europa Lander concept, showing higher utility achievement compared to baseline approaches.

1 INTRODUCTION

When integrating AI planning into robotic applications, planners are consistently challenged by variation in execution and uncertainty in the quality of our environment models. In space-based applications, this is especially challenging because the environment is largely unknown, reducing the quality of our a priori models of the world. To address these problems, we describe an integrated approach to planning and execution in an unknown, unpredictable environment. We use a hierarchical task network (HTN) that defines activities and their associated dependencies, and we create plans using this mission representation. To drive planning goals, we assign utility to tasks that complete those goals, and generate a plan that maximizes utility while obeying each task's constraints. At execution time, we use flexible execution and re-planning to react to uncertainty and variation, and to compensate for an approximate world model. Finally, we accept model parameter updates (such as expected task energy usage and estimated task utility) during execution time to incorporate execution feedback and gained knowledge and use these updated models to optimize plan quality.

We examine the problem in the context of a proposed mission concept to perform in-situ analysis of samples from the surface of the Jovian moon Europa [1]. Unlike prior NASA missions, a priori domain knowledge is severely limited and uncertain, and communication with Earth is limited by long blackout periods (about 42 hours out of every 84 hours). Consequently, a successful mission requires a planning and execution framework that is highly efficient, robust to unprecedented levels of uncertainty, and still capable of maximizing its overall utility. On the other hand, the Europa Lander concept has a fairly rigid definition of what actions the lander must perform in order to produce utility. Our planning algorithm leverages this domain-specific knowledge by making use of a hierarchical task network and using heuristic-guided search to examine various task combinations to maximize utility. The ultimate goal for a Europa Lander would be to analyze surface material and communicate the resulting data products back to Earth. To reward accomplishment of these goals, we assign utility to tasks such as sample excavation and seismographic data collection, but do not receive this utility until the lander communicates the data down to Earth. In the HTN framework, this means that tasks in a hierarchy produce utility only if the full hierarchy is executed.

We integrate planning and execution in order to better react to environmental variation, moving away from fixed time and energy budgets to generate less conservative, more successful plans. To do so, we use MEXEC, an integrated planner and executive first built for NASA's Europa Clipper mission [2]. To account for approximate world models, our framework replans on a periodic basis to recalibrate its plans with reality. Finally, we integrate a module that continually estimates task parameters and goal utilities during execution in response to new information. We update the priors on the relevant model parameters accordingly, and subsequently replan to produce a more accurate plan. This further improves plan quality and utility gain.

We present this planning and execution framework in a simulated Europa-like mission and compare it to three baseline approaches similar to those used in prior missions: a static plan [3], flexible execution with no replanning, and flexible execution with replanning but without online parameter update [4]. We explore the value of flexible execution, replanning, and online model reparameterization, and examine their effect on utility in these scenarios. We present empirical results of the efficacy of each technique over the others, and discuss further techniques and implementations that may continue to improve on these baselines into the future.

2 PROBLEM DESCRIPTION

The primary goal of the Europa Lander mission concept is to excavate and sample the surface, analyze the sampled material for signs of biosignatures, and communicate that data back to Earth [1]. Additionally, there are secondary objectives to take panoramic imagery of the European surface and collect seismographic data. Lander operations are generally limited to the accomplishment of these two overarching goals. This provides significant structure to the problem, since the concept mission clearly defines the sequence of actions required to achieve these goals.

As a minimum requirement, the lander should excavate a trench in the European surface, collect three samples from that site, analyze those samples, and return that data to Earth. The basic requirements of a mission would require only a single site to be excavated. However, there is value in excavating additional sites, because the material at different sites may possess different properties. In addition, the lander may choose to resample the same location, for example, in order to verify the discovery of a biosignature at that location. In the baseline mission concept, all three of the lander's samples are chosen from the same target. Note that after the first site is excavated, no further excavations are needed to sample from that trench; all three sampling activities can share a single excavation site. After excavation and sample collection, samples must be transferred into scientific instruments that analyze the material and produce data products. Then, for a mission to achieve any actual utility, those data products must be communicated back to Earth. Because communication is difficult and energy intensive, the lander may choose to compress data lossily if the expected utility of this action is higher.

In addition to sampling tasks, the lander may engage in seismographic data collection and period panoramic imagery tasks. These are considered lesser goals, with lower utility associated with their completion. As

such, the data products that these tasks generate are considered to have lower value. However, these tasks also involve no surface interaction, and have less uncertainty associated with them as a result.

It is important to note that utility is only achieved when data is downlinked back to Earth. This is true for both the sampling and seismograph/panorama tasks. Some excavation sites or sampling targets may provide more utility than others if, for example, one of those targets has a positive biosignature and the other does not. However, regardless of the quality of the material that the lander samples, no utility is achieved unless that data is communicated. This dynamic means that while potential utility is generated during the sampling and analysis phases, it is only realized by completing relevant communication tasks.

The Europa Lander mission concept is also constrained by a finite battery that cannot be recharged. Battery life is a depletable resource, and the lander must use its energy as efficiently as possible. Each task saps energy from the battery, and our algorithm must plan accordingly to maximize utility in face of this constraint. In addition to this challenge, the surface characteristics of Europa are uncertain, and any prior mission model that is generated before landing is sure to have inaccuracies. In particular, the energy consumption of the excavation and sample collection tasks is largely unknown. There is also significant variation in the utility of any given sample, since the value of sampling a given target on Europa depends on whether the material is scientifically interesting, e.g. whether a biosignature is present.

3 APPROACH

3.1 Problem Model

We model this problem using a hierarchical task network (HTN) to compile the domain-specific knowledge of the dependency structure into the task network. HTNs have been used successfully in industrial and other real-world applications to improve the tractability of planning problems in systems such as SHOP2 [5] and SHOP3 [6]. In an HTN, hierarchical tasks are decomposed to a set of subtasks. We refer to the higher-level tasks as “parent tasks”, and refer to their children as “subtasks”. Parent tasks may decompose into a number of different sets of subtasks; we refer to each of these sets as a potential “decomposition” of that parent task. Finally, we refer to tasks with no decompositions as “primitive tasks”. These primitive tasks represent tasks that the lander can be directly commanded to perform.

/Decompositions provide a number of benefits to our planning approach, significantly reducing plan search space. In addition, we can treat all subtasks of a parent task as a singular block for planning purposes. The lander only achieves utility after completing an entire sequence of sample, analyze, communicate. Decompositions allow us to treat “sample, analyze, communicate” as a single unit and schedule them accordingly. Thus, our model intrinsically biases the lander against planning to sample without a corresponding communication task. This may not always be optimal, if for example, excavation and sampling is cheap and communication is very expensive. However, for our problem, energy use is dominated by the excavation and sampling tasks, and the decomposition paradigm effectively encodes this domain-specific knowledge into our planning routine.

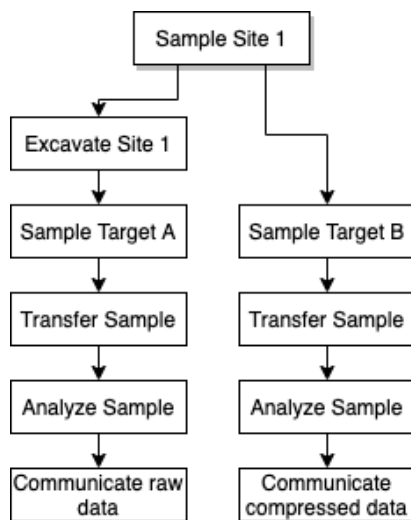


Figure 1: Example set of two possible decompositions for a “Sample Site” task.

There are three main parent task types in our mission model. The first is a Preamble, which consists of post-landing initialization and other one-time initialization tasks. Second are sampling tasks. These consist of excavation, sample collection, transfer, analysis, and communication tasks. Excavation can take place at one of two excavation sites, and may be skipped if an excavation has previously occurred for the specified site. For collection tasks, the lander may choose between four collection targets: two for each excavation site. It may revisit a target that has already been sampled, still obtaining utility for a repeat sample. Then, for communication tasks, the lander may choose to either communicate raw data or compressed data. Finally, there are Seismograph/Panorama tasks, which

consist of seismographic data collection, panoramic image collection, and communication of that data.

In our problem, we assign utility primarily to two activities: sampling and communication. We assign utility to sampling tasks in order to differentiate between sites that may be more or less interesting, depending on the scientific value of the site. Communication utility is larger, and remains constant. For the communication tasks, we assign higher utility and cost to tasks that communicate raw data, compared to those that communicate compressed data. This simulates a Pareto optimal “menu” of communication options. The combination of sampling and communication utilities represents the overall utility of a parent sampling task. Seismograph/panorama utility is driven solely by communication utility.

3.2 Planning Algorithm

Our planning algorithm uses the HTN model of the Europa Lander problem to build a search graph, with nodes holding partial plans and edges holding task decompositions. We perform a heuristic-guided branch and bound search on this graph and select the best plan explored. The algorithm consists of four phases: pre-processing, initialization, exploration, and plan selection.

First, a pre-processing step flattens task decompositions into a single layer, such that parent tasks decompose into a chain consisting only of primitive, non-hierarchical subtasks. This allows us to assign utility and energy cost directly to each decomposition, because its breakdown into disparate subtasks has already been performed. Then, each decomposition's utility is the sum of each of its subtasks' utility. The same is true for energy cost. This step is performed once per domain model, offline. Preprocessing has exponential runtime in the worst case, and future work may require additional search in decomposing tasks as well as planning them.

Our search graph consists of nodes containing partial plans and their associated energy cost and utility. A node's cost is simply the sum of cost of each task scheduled in the node's plan; the same goes for utility, though future work may take joint utility into account. In the initialization phase, the algorithm creates a single node containing an empty plan, with utility and cost 0. Then, it iterates through all task decompositions created in the pre-processing phase in order to generate the set of edges that may be followed from a given node. To finish the initialization phase, the algorithm populates an exploration queue with (node, edge) pairs, pairing the singular initial node with all

edges in the collection. At the end of the initialization phase, then, the exploration queue consists of all task decompositions paired with the empty plan.

In the exploration phase, the planner pops the top of the exploration queue to get (P, T) , where P is a partial plan, and T is the list of primitive subtasks comprising a task decomposition. It then attempts to schedule all tasks in T given the state of the world produced by following the plan P . If the tasks cannot be scheduled, it moves on to the next exploration queue item. If the tasks can be scheduled, i.e. their preconditions are met and their impacts do not produce any conflicts, a new graph node is created. This node contains a new plan P' , the resulting plan after adding the tasks in T to P .

After creating this plan node, the planner iterates through the edge collection again, pairing the new plan with all possible tasks. In this iteration, it ignores tasks that have already been scheduled in the plan, so as to avoid duplicates. The algorithm also filters these pairs to ensure that the total cost $P.cost + T.cost < M$, where M is the max energy cost allowed (equal to the current battery charge of the lander). This bounds our search, and we further bound the algorithm's search by limiting the number of exploration candidates examined. Note however that this bound maintains optimality if we allow the algorithm to expand the entire space. After filtering, these pairs are added to the exploration queue, and the next queue item is examined. The exploration queue is a priority queue, with (plan, decomposition) pairs ordered by a heuristic value to improve search results. Given a plan, decomposition pair (P, T) we assign the heuristic value $h(P, T) = P.utility + T.utility/T.cost$. Finally, in the plan selection phase, the algorithm iterates through all candidate plan nodes, selecting the plan with the highest utility. Ties are broken according to energy cost, where a lower energy cost is preferred.

3.3 Execution

Planning and execution are integrated in our approach, in order to respond to variation and therefore better optimize overall utility achieved. We use MEXEC for flexible execution, which allows us to make small plan modifications to task start times and resource consumption without failing [2]. We maintain task models that have impacts on certain resources, e.g. each task has some negative impact on an energy resource, representing its energy usage. As is common in robotics, variation during execution time may prove these models to be incorrect. Thus, during execution, we update our resource timelines to match the values measured.

In addition to updating the raw value of modeled resources, we also update task impact models with the information gained during execution. This is especially important in the sampling tasks. There exists significant uncertainty in our model of interactions with the European surface, so we may discover that collecting a sample is much more difficult than previously expected. In this case, we rely on updating our model of the sample collection task to better represent what has been discovered. By doing so, we are able to replan with a better understanding of the task, thereby creating plans that are more likely to succeed. In this paper, we explore only one online parameter update policy, where the latest energy use value is used to update the task model, i.e. our prior estimate is ignored after gaining some posterior knowledge.

We also update our utility estimation for certain tasks. Our task network has a prior estimate on the utility of collecting surface samples that is uniform across all targets. At execution time, the information discovered may drastically alter this estimation. For example, after sample analysis, the lander may discover a biosignature at target A, but nothing at target B. Our framework would then update its task models accordingly, rewarding sample collection tasks at target A with much higher utility than tasks at target B. Online utility update thus allows us to redefine our goals and the value of each goal according to the information obtained at execution time.

Tying all of this together is periodic replanning on a fixed cadence. This allows us make use of online state updates, model updates, and utility reassignments during execution time, thereby creating plans that achieve higher overall utility. Our framework measures the value of each resource being modeled, and assigns that value to the given resource in the planning model. Then, when replanning, the planner uses the actual, measured value of the state, rather than the previous predicted value. This allows us to update our goals according to what is realistically possible given the current state measurements of the system. In future work, we may replan in response to events such as the detection of significant error in resource modeling [7]. This would minimize replanning, such that it would occur only when necessary.

4 RESULTS

We examine the performance of four potential planning and execution frameworks in a simulated Europa Lander problem. As a baseline, we examine the performance of a static plan, executed with fixed time points and energy impacts. We refer to this strategy as

/"none", as it uses none of the flexible execution techniques we describe. While this strategy is extremely basic, it is also the baseline for many NASA missions, including the Curiosity rover [2][8]. With this strategy, we generate a static plan and attempt to execute it, without any flexibility. Because we lack flexibility, the static plan must contain margin such that tasks are unlikely to overrun their allotted resource budgets. If any overruns do occur, we terminate the execution of the plan.

We refer to the second strategy as "flexible". With this strategy, we allow flexible execution of a plan using MEXEC, but do not replan during execution. Because execution is flexible, we can reduce the amount of margin in our task models, and rely on our executive to handle resource overruns accordingly. In a real application, this would be analogous to placing an executive on a spacecraft, but only allowing it to run a plan generated on the ground.

The third strategy is "replan", where we allow for replanning during execution, but do not perform any model parameter or utility updates. At a periodic cadence, we run the planning algorithm using updated resource values, but we do not incorporate any information gained about task model parameters or utility. This baseline is comparable to expected operations on the Mars 2020 mission [4].

Finally, the "model_update" strategy incorporates all techniques described in the approach, executing plans flexibly, replanning on a regular cadence, and updating task model parameters and goal utilities online. Specifically, we estimate and update parameters for the sample collection task. This task is particularly important because it uses significant energy and is repeated many times in the plan. In addition, as a task that interacts with the European surface, our model of the task is likely to be extremely uncertain.

To evaluate the efficacy of these approaches, we simulate execution on a Europa Lander problem. In this simulation, we vary the energy usage of each task, drawing from a normal distribution centered around the prior estimate in the input task network. In addition to this variation in energy, we vary the utility of different sampling targets. We draw this value from a high-variance normal distribution, again centered around the prior estimate in the task network. Finally, to simulate uncertainty in our model, for the sample collection task, we draw energy use from a normal distribution centered around a mean that does not match the task network input. Instead, the mean is drawn from another normal distribution, centered around the input.




Figure 2. Average utility achieved by the four planning and execution strategies applied to simulated Europa Lander scenarios.

This means that the energy distribution expected in the world model may not match the true distribution.

The results of our experiment are shown in Fig. 2. Each strategy was simulated for 50 runs, recording the average utility achieved by the execution trace. We find that the "none" strategy's static plan is severely limited by the margin required to ensure that the plan can be executed without resource overruns. This margin is two standard deviations above mean, and limits our plan to only sampling a single time. Even with this margin, we find that resource overruns still occur, which result in no utility achieved. The static approach achieves a median utility of 1585 and mean of 1325.

Allowing flexible execution with the "flexible" strategy improves this significantly, bumping median utility gain to 3911, and mean utility to 4195. This benefit comes primarily in the removal of the margins. With an onboard executive, our plans can expand to the full set of 3 samples, which greatly increases the ceiling for our utility gain. Fig. 3 contrasts the plans of the "none" and "flexible" strategies at a high level. These plans were selected from execution traces of the experiment. While flexible execution is very valuable, because we do not replan, we find that the utility achieved by this strategy varies widely. Plans are fairly brittle to unexpected energy loss, often resulting in missing the final communication step in the plan. This proves disastrous for utility gain, because the work required to generate the final data product is essentially wasted. Thus, we see a wide spread of overall utility in this strategy, with two main clusters:

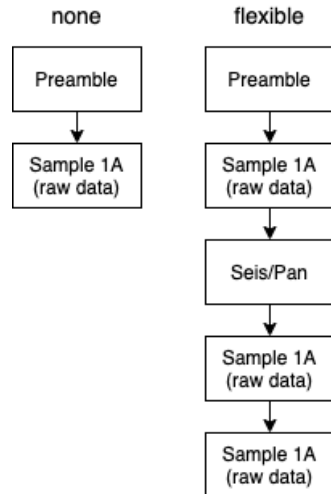


Figure 3. Sample high-level plans for the none and flexible planning/execution strategies.

one for traces where all three samples are successful, and one for traces where one sample fails.

With the “replan” strategy, we allow replanning, which significantly reduces the brittleness of our plans to variation. We find that this reduces the variance of utility achieved, but does not significantly affect the overall utility gained (median 4185, mean 4027). Fig. 4 shows a sample plan with this strategy. Because we react to variation in energy, the planning algorithm begins to favor compressed data as a way to save on energy while still achieving sampling goals. However, this also cuts down on potential utility. Because the flexible strategy is so optimistic, it is bimodal in utility gain, either achieving close to optimal, or falling far short. The replan strategy tends toward a middle path, with a lower variance of moderate utility plans. This variance reduction is useful for the Europa Lander problem, since the cost of failure is significant.

Finally, the “model_update” strategy results in a median utility gain of 4372 and mean of 4327. The most significant difference in the full strategy is the online re-weighting and re-parameterization of different sampling tasks. Depending on what is discovered at a given sample target, we update the utility of the sampling action corresponding to that target. This motivates the lander to sample in other locations if a particular target is found to be less valuable or more costly than expected. Because the planner can take advantage of this newfound knowledge, it is able to harvest greater utility from sample targets. In our sample plan, we find that this case occurs in practice: After discovering that sampling target A is less interesting or more difficult than originally predicted, the lander chooses

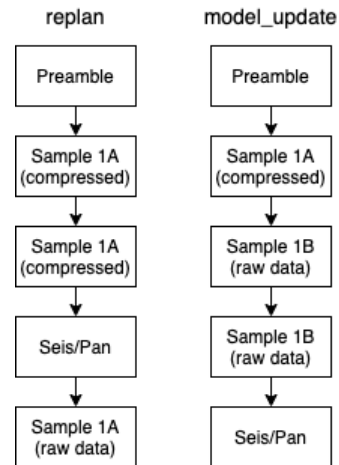


Figure 4. Sample high-level plans for the replan and model_update planning/execution strategies

to swap targets to target B. It is thus able to make better use of energy and more effectively harvest utility, compared to methods that do not estimate parameters online.

5 RELATED WORK

Onboard planning and execution are of great interest to the space domain. The *Remote Agent* was an architecture for onboard planning and execution addressing remote autonomous operation with deadlines, resource constraints, and concurrent activities [9]. The Remote Agent flew for 48h in 1999 on the Deep Space One spacecraft using a batch planner that took hours on a RAD6000 CPU to generate a temporally flexible plan that was then used by a reactive executive controller [10] to provide robust plan execution. The planner used a refinement search paradigm [11] to construct a temporally flexible plan but did not consider utility in plan generation and did not perform continuous replanning due to the computational expense and long planning time (indeed the replans were scheduled in the prior plan).

The Earth Observing One (EO-1) spacecraft [12], which flew for over 12 years from 2004-2017, was designed specifically to react to dynamic scientific events. Planning was performed by the CASPER planning software [13], which took on the order of 10s of minutes to replan but did not produce temporally flexible plans. To address this, the onboard executive (SCL) was able to flexibly interpret the *execution* of a plan to handle minor execution runtime variations. The flight and ground planners [14] both used a domain specific search algorithm that enforced a strict priority model over observations for limited model of

utility. This scenario is similar to that proposed in this paper, in which the lander must react to dynamic events and observations in order to maximize its utility, while still adhering to both mission and spacecraft constraints. Recently, the Intelligent Payload Experiment (IPEX) also successfully used the CASPER planning software to achieve its mission objective, further validating the efficacy of using onboard replanning to handle dynamic events and observations during operation even when the plans are not temporally flexible [15].

The M2020 Perseverance rover also plans to fly an onboard planner [4] to reduce lost productivity from following fixed time conservative plans [8]. Like the planning approach we propose in this paper, the M2020 planning architecture also relies on rescheduling and flexible execution [7], ground-based compilation [16], heuristics [17], and very limited handling of planning contingencies [18]. However, many characteristics of the M2020 mission are fundamentally different from the mission concept we consider here, such as the lack of reliable a priori model parameters, the inability to recharge the battery, and the long communications blackout time windows incentivizing greater mission autonomy.

6 FUTURE WORK

In our planning, we consider only one dimension of decision theory: utility. While we react to uncertainty at execution time, we do not take this into account during the planning phase. A more sophisticated planner would explicitly integrate probability into plan generation, maximizing expected utility rather than assuming resource impacts are constant and correct. For example, excavation tasks involve risk; task failure could result in significant energy loss or damage to the lander. Reasoning about exogenous events such as these would improve utility achievement by potentially avoiding such risks, or even seeking them out later in the mission when failure is less impactful.

In addition, considering depletable resource usage probabilistically would allow optimization of plans to avoid significant utility loss if that resource is depleted. For Europa Lander, energy use is a significant constraint on all tasks. By reasoning about the probability of running out of energy, we can improve our plans and avoid these cases when possible. One simple approach would be to discount a task's utility according to the energy remaining at the time it is scheduled. This is analogous to the intuitive approach of discounting a task's utility according to its distance into the future. Instead of tying utility discounts to time,

however, we tie it to energy, because this is the most significant task constraint in the Europa Lander problem.

A more sophisticated approach would examine plan prefixes and choose the one that best balances potential utility and risk. Here, we would generate a set of plan prefixes that plan until some fixed time or energy point is reached, e.g. the start of the next planning cycle. Given this plan prefix, we would generate plan postfixes given various energy levels, drawn from a distribution representing the battery remaining. Using a sampling approach, we could arbitrarily refine our estimate of the plan prefix's quality given various energy levels. Then, we would pick the plan prefix that maximizes the overall expected utility of the entire plan. Adding probabilistic reasoning into our planning algorithm is likely to improve plan quality and thus improve the overall performance of our planning and execution framework.

Acknowledgement

Pre-decisional information – for planning and discussion purposes only. This work was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. © 2020. California Institute of Technology. Government sponsorship acknowledged.

References

- [1] Hand, K. (2017). *Report of the Europa Lander science definition team*. National Aeronautics and Space Administration.
- [2] Vandi Verma, Dan Gaines, Gregg Rabideau, Steve Schaffer, & Rajeev Joshi (2017). *Autonomous Science Restart for the Planned Europa Mission with Lightweight Planning and Execution*. In *International Workshop on Planning and Scheduling for Space (IWSS 2017)*.
- [3] Gaines, D., Anderson, R., Doran, G., Huffman, W., Justice, H., Mackey, R., Rabideau, G., Vasavada, A., Verma, V., Estlin, T., & others (2016). *Productivity challenges for mars rover operations*. In *Proceedings of 4th Workshop on Planning and Robotics (PlanRob)* (pp. 115–125).
- [4] Gregg Rabideau, & Ed Benowitz (2017). *Prototyping an Onboard Scheduler for the Mars 2020 Rover*. In *International Workshop on Planning and Scheduling for Space (IWSS 2017)*.
- [5] Nau, D., Au, T., Ilghami, O., Kuter, U., Murdock, J., Wu, D., & Yaman, F. (2003). *SHOP2: An HTN Planning System*. In *Journal of Artificial Intelligence*

Research, 20, 379–404.

[6] Goldman, U. (2019). *Hierarchical Task Network Planning in Common Lisp: the case of SHOP3*. In *Proceedings of the 12th European Lisp Symposium* (pp. 73–80). Zenodo.

[7] W. Chi, S. Chien, J. Agrawal, G. Rabideau, E. Benowitz, D. Gaines, E. Fosse, S. Kuhn, & J. Biehl (2018). *Embedding a Scheduler in Execution for a Planetary Rover*. In *International Conference on Automated Planning and Scheduling (ICAPS 2018)*.

[8] Gaines, D.; Doran, G.; Justice, H.; Rabideau, G.; Schaffer, S.; Verma, V.; Wagstaff, K.; Vasavada, A.; Huffman, W.; Anderson, R.; et al. 2016b. *Productivity challenges for mars rover operations: A case study of mars science laboratory operations*. Technical report, Technical Report D-97908, Jet Propulsion Laboratory.

[9] Muscettola, N., Nayak, P., Pell, B., & Williams, B. (1998). *Remote agent: To boldly go where no AI system has gone before*. *Artificial intelligence (AIJ)*, 103(1-2), 5–47.

[10] Pell, B., Gat, E., Keesing, R., Muscettola, N., & Smith, B. (1997). *Robust periodic planning and execution for autonomous spacecraft*. In *IJCAI* (pp. 1234–1239).

[11] Jonsson, A. K.; Morris, P. H.; Muscettola, N.; Rajan, K.; and Smith, B. D. 2000. *Planning in interplanetary space: Theory and practice*. In *AIPS*, 177–186.

[12] Chien, S., Sherwood, R., Tran, D., Cichy, B., Rabideau, G., Castano, R., Davis, A., Mandl, D., Frye, S., Trout, B., & others (2005). *Using autonomy flight software to improve science return on Earth Observing One*. *Journal of Aerospace Computing, Information, and Communication*, 2(4), 196–216.

[13] Chien, S., Knight, R., Stechert, A., Sherwood, R., & Rabideau, G. (2000). *Using Iterative Repair to Improve the Responsiveness of Planning and Scheduling*. In *International Conference on AI Planning and Scheduling (AIPS)* (pp. 300–307).

[14] Chien, S., Tran, D., Rabideau, G., Schaffer, S., Mandl, D., & Frye, S. (2010). *Timeline-based space operations scheduling with external constraints*. In *Twentieth International Conference on Automated Planning and Scheduling*.

[15] Chien, S., Doubleday, J., Thompson, D., Wagstaff, K., Bellardo, J., Francis, C., Baumgarten, E., Williams, A., Yee, E., Stanton, E., & others (2017). *Onboard autonomy on the intelligent payload experiment cubesat mission*. *Journal of Aerospace Information Systems*, 14(6), 307–315.

[16] W. Chi, J. Agrawal, S. Chien, E. Fosse, & U. Guduri (2019). *Optimizing Parameters for Uncertain Execution and Rescheduling Robustness*.

In *International Conference on Automated Planning and Scheduling (ICAPS 2019)*.

[17] W. Chi, S. Chien, & J. Agrawal (2020). *Scheduling with Complex Consumptive Resources for a Planetary Rover*. In *International Conference on Automated Planning and Scheduling (ICAPS 2020)*.

[18] J. Agrawal, W. Chi, S. Chien, G. Rabideau, S. Kuhn, & D. Gaines (2019). *Enabling Limited Resource-Bounded Disjunction in Scheduling*. In *11th International Workshop on Planning and Scheduling for Space (IWPSS 2019)* (pp. 7–15).