

Analyzing the Efficacy of Flexible Execution, Replanning, and Plan Optimization for a Planetary Lander

Daniel Wang, Joseph A. Russino, Connor Basich, and Steve Chien

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, California 91109

Abstract

Plan execution in unknown environments poses a number of challenges: uncertainty in domain modeling, stochasticity at execution time, and the presence of exogenous events. These challenges motivate an integrated approach to planning and execution that is able to respond intelligently to variation. We examine this problem in the context of the Europa Lander mission concept, and evaluate a planning and execution framework that responds to feedback and task failure using two techniques: flexible execution and replanning with plan optimization. We develop a theoretical framework to estimate gains from these techniques, and we compare these predictions to empirical results generated in simulation. These results indicate that an integrated approach to planning and execution leveraging flexible execution, replanning, and utility maximization shows significant promise for future tightly-constrained space missions that must address significant uncertainty.

Introduction

AI planning for robotic applications often must address variation in execution and uncertainty in the quality of environment models. In space-based applications, this can be especially challenging when 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. First, we define a theoretical framework to examine the value of two integrated planning and execution techniques: flexible execution and replanning with plan optimization. We discuss this framework in the context of the Europa Lander mission concept. Finally, we compare the predictions of the model to empirical results in a Europa-like simulation environment.

The primary empirical context of our model is a mission concept to perform *in situ* analysis of samples from the surface of the Jovian moon Europa (Hand 2017). Unlike prior NASA missions, *a priori* domain knowledge is severely limited and uncertain, and communication with Earth is limited by long blackout periods (over 42 hours out of every 84 hours). Consequently, a successful mission requires

a planning and execution framework that can operate autonomously for extended periods of time, is robust to unprecedented levels of uncertainty, and is still capable of maximizing its overall utility. Additionally, because of the harsh radiation environment at Europa, mission lifetime and on-board computing are severely limited¹.

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 (HTN) 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 the overwhelming majority of the mission utility is not awarded until the lander communicates the data down to Earth. In the HTN framework, this means that tasks in a hierarchy produce very little utility until the full hierarchy is executed.

For our empirical evaluation, we base our planning system on MEXEC, an integrated planner and executive originally built for NASA's Europa Clipper mission (Verma et al. 2017). We compare four approaches to planning on the Europa Lander problem similar to those used in prior missions: a static plan without failure recovery mechanisms, a static plan with ground input for failure recovery (similar to current Mars Rover operations) (Gaines and et al 2016), flexible execution without replanning, and flexible execution with replanning optimization. We explore the value of onboard autonomy: flexible execution and replanning with plan optimization, and examine these techniques' effects on utility in these scenarios. We demonstrate that, true to our model's prediction, each technique shows significant improvement in utility achievement in the Europa Lander domain.

¹As a point of reference, the RAD750 processor used by the Mars 2020 rover has measured performance in the 200-300 MIPS range. In comparison, a 2016 Intel Core i7 measured over 300,000 MIPS, or over 1000 times faster. Furthermore, the Mars 2020 on-board scheduler (Agrawal et al. 2021b) is only allocated a portion of the computing cycles onboard the RAD750 resulting computation *several thousand times* slower than a typical laptop.

Domain 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 (Hand 2017). 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. Figure 1 displays the strong dependency structure inherent to the Europa Lander concept mission. In order to sample, the lander needs to have excavated a trench; in order to analyze, the lander needs to have collected a sample; etc.

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. On the other hand, the lander may choose to resample the same location, for example, in order to verify the discovery of a biosignature. 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.

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 primary utility is only achieved when data is downlinked back to Earth². 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

²While context imagery while preparing sampling locations has some science value, the primary science value is from analysis of the samples.

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.

Approach

We design our planning system to respond intelligently to stochasticity at execution time, since we expect this to be a significant factor in our domain. Planning and execution are integrated in our approach, in order to respond to variation and therefore better optimize overall utility achieved. We achieve this integration through the use of two techniques: flexible execution and replanning with plan optimization.

Flexible Execution

Flexible execution is a lightweight rescheduling algorithm that runs at a much higher cadence than the planner. This algorithm has two main properties: (1) it is much less computationally demanding than replanning as it does not search, and (2) it is less capable than replanning. Flexible execution allows the system to handle less-severe unexpected events without incurring the cost of replanning. Previous NASA missions have made heavy use of flexible execution, such as the Mars 2020 Perseverance rover (Agrawal et al. 2021a). Our implementation differs in focus, emphasizing responses to adverse events.

In our system, flexible execution consists of two major components. The first is *task push*. If a task's preconditions are not met, before failing the task, we allow it to wait for some amount of time for this inconsistency to resolve. Such a situation might occur, for example, if required preceding activities are delayed or run long. The executive checks the task's preconditions and delays dispatch until either the conditions have been met, or the task's wait timeout has been exceeded.

The second component of flexible execution is *automated retry*. If a task completes with a failure code, flexible execution can immediately re-schedule the task if its preconditions are still met (and plan updated with new predicted end time of the task and resource usage), avoiding replanning cost and delay.

In the context of the Europa Lander domain, flexible execution offers significant value because many robotic tasks such as trenching and sample acquisition can vary significantly in duration and hence resource consumption. Flexible execution handles this variation without disrupting the execution flow.

Replanning with Plan Optimization

For more complex execution variations, we turn to replanning during execution. Replanning uses search to construct

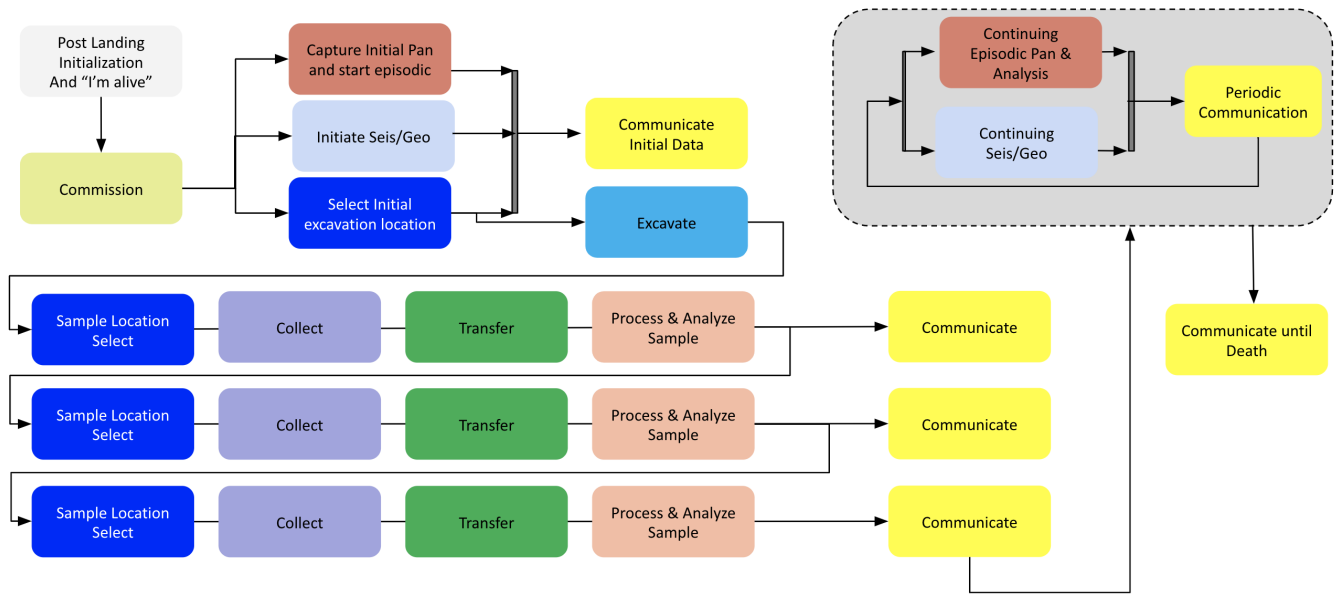


Figure 1: A task network for the Europa Lander mission concept. The diagram represents a potential execution trace of the mission that would fulfill baseline requirements.

the plan based on the current state. The current state includes: state and resource values (accounting for failed activities, and resource under/over runs) as well as the current time (accounting for execution time variation). Additionally, utility estimation, predicted duration, and predicted resource expenditures for future tasks may be updated (e.g. imaging may indicate that a sampling site now looks more promising and/or may take longer/shorter to excavate or sample). Re-planning enables incorporation of all of this new information into decision-making.

Theoretical Framework for Analysis

We define our planning problem as follows. We provide our planner with a set of tasks $T = \{t_1, \dots, t_n\}$. Each task is represented by a tuple $t_k = \{c_k, u_k, d_k, P, I\}$ where:

- c_k represents the task's cost.
- u_k represents the task's utility.
- d_k represents the task's nominal duration.
- P is the set of the task's preconditions. These may be based on resource values, or on the execution state of dependency tasks.
- I is the set of its impacts on resource timelines.

This matches the timeline representation of execution state used by (Verma et al. 2017). For our problem, we assume that we have a fixed cost budget b . In the Europa Lander domain, this budget represents the non-rechargeable battery, with each task using up some amount of that battery's energy. We wish to maximize utility by scheduling tasks subject to the following constraints:

- For all tasks, all preconditions are valid.
- For all tasks, all impacts are valid.

- The sum of all task costs does not exceed b .

In our framework, we examine four planning and execution strategies of increasing onboard autonomy: static, ground, flexible execution (FE), and replan. In static, a plan is generated before execution time, then executed without change. No failure responses are available, so the first task failure results in the remainder of the plan not executing. In ground, we introduce a mechanism for failure resolution: waiting for ground input. We assume that ground input is able to resolve all failures. The plan generated on the ground, and any task failures result in a halt to plan execution, ground fixes the problem, then execution resumes, albeit incurring considerable cost. In the FE strategy, we allow flexible execution of our plans, which can resolve some but not all failures, with all other failures handled by waiting for ground input. Finally, in the replan strategy, first FE is applied to resolve failures, if FE cannot resolve the failure replanning is applied, if replanning cannot resolve the failure, the ground strategy is applied. Replanning can therefore serve dual purposes: resolving task failures, and replanning to increase plan utility.

Given this context, we predict the overall utility achievement of a plan using an estimate of utility per unit cost u_{avg} . If all tasks in a plan always succeed, our expected utility for a plan would be bu_{avg} .

To factor in task failure, we assume that tasks fail with some probability $P(\text{fail})$, and we assume that task failures follow a Poisson distribution in that each task fail with a fixed failure rate ³.

We assume that some subset of these failures can be resolved with FE, a strictly larger subset can be resolved using replanning and that all failure modes can be resolved via

³Extending the analysis to more realistic plan structure, failure patterns, and interactions with exogenous events are future work.

waiting for ground input. Then, if $P(\text{ground})$ denotes the probability that a failure is resolvable by ground (but not replan or FE) and where $P(\text{replan})$ is the probability that replan works but FE does not, failure handling probabilities can be described as:

$$P(\text{fail}) = P(\text{ground}) + P(\text{replan}) + P(\text{FE}) \quad (1)$$

We first analyze the static strategy. Assuming the plan is a linear sequence of activities and execution terminates execution with the first failure, the expected number of successfully executed activities is $1/P(\text{fail})$ but bounded by the number of tasks n . Therefore the static expected utility is the lower of the "no failure case" (at left below) or (at right) the expected utility per task $u_{avg} \cdot c_{avg}$ times the number of expected activities executed as indicated below:

$$U(S_s) = \min \left(u_{avg} \cdot b, \frac{u_{avg} \cdot c_{avg}}{P(\text{fail})} \right) \quad (2)$$

where c_{avg} denotes the average cost of each task.

In the ground strategy, the only error response for all failures is to "go to ground" to resolve which always allows plan execution to continue. If our plan has n tasks, this occurs $nP(\text{fail})$ times costing c_g energy each occurrence. The utility achievement of this "ground" strategy is:

$$U(S_g) = u_{avg} (b - P(\text{fail})nc_g) \quad (3)$$

In the FE strategy, we introduce flexible execution and assume that some subset of task failures can be resolved with this feature. The probability of a task failing in this way is $P(\text{FE})$ and note that $P(\text{FE})$ is a subset of the total failure cases $P(\text{fail})$. We assume that flexible execution has a negligible cost. Then, the utility achievement of plan execution using this strategy is:

$$U(S_f) = u_{avg} (b - (P(\text{fail}) - P(\text{FE}))nc_g) \quad (4)$$

Finally, we consider the replan strategy, which incorporates flexible execution and replanning with plan optimization. Unlike flexible execution, replanning incurs some non-negligible cost c_r . We assume that, like flexible execution, replanning is able to resolve some subset of task failures. We denote the probability that a given task fails in a way that can be resolved via replanning, but not flexible execution, as $P(\text{replan})$.

Failures are resolved by the least costly resolution mechanism. Thus, when a task fails, our system attempts to resolve it by flexible execution, if possible, falling back to replanning and ground intervention in sequence. To model plan optimization, we provide our planning system with opportunities to discover utility at certain points during execution. We denote the number of such opportunities as d , and the expected additional utility discovered as u_d . Then,

$$U(S_r) = du_d + u_{avg} (b - n(P(\text{ground})c_g + P(\text{replan})c_r)) \quad (5)$$

Planning Approach

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

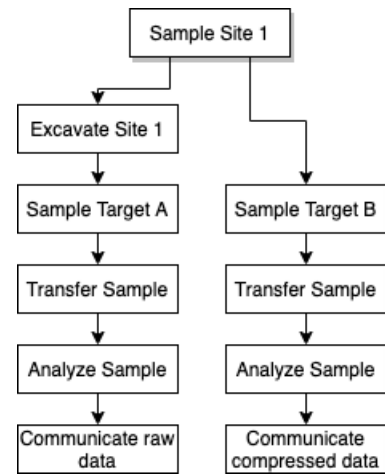


Figure 2: Two possible decompositions of a single parent "Sample Site 1". In the left decomposition, the lander excavates the site, samples target A, and communicates raw data. In the right decomposition, the lander skips excavation, samples site B, and communicates compressed data. Both achieve the same goal of sampling site 1.

used successfully in industrial and other real-world applications to improve the tractability of planning problems in systems such as SHOP2 (Nau et al. 2003) and SHOP3 (Goldman and Kuter 2019). 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 primarily 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.

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. Both of these task models are assigned a numeric value representing their utility, which can be updated online by the planning and execution system if knowledge at execution time alters the expected utility of a given action. Utility for these tasks is achieved only after their full decomposition has been successfully executed. Thus, for sampling utility to be achieved, a corresponding communication step must successfully complete.

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.

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 expected 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. While this pre-processing has exponential runtime in the worst case, for our Europa Lander (as well as most space applications we have seen) the majority of the search occurs in scheduling the expanded tasks not action selection so this preprocessing is tractable (managing this complexity is an area for future work).

Our search tree consists of nodes containing partial plans and their associated energy cost and utility. A node’s cost is the sum of the costs of each task in the node’s partial plan; and likewise the node utility is the sum of the utilities of the (sub) tasks (task utility that is dependent on other tasks is a topic for future work). In the initialization phase, the algo-

Algorithm 1: Europa Lander Planning

```

Input: A list of tasks to schedule  $T$ 
Output: A plan of scheduled tasks  $P$ 
/* initialize exploration queue */
node_collection = [];
add (plan=[], utility=0, cost=0) to node_collection;
edge_collection = [];
for  $d$  in task.decompositions do
    new_edge = ( $d$ ,  $d$ .utility,  $d$ .cost);
    add new_edge to edge_collection;
end
explore_q = [];
for edge in edge_collection do
    add (node_collection[0], edge) to explore_q;
end
/* search exploration queue */
num_explored = 0;
while num_explored below exploration bound do
    num_explored++;
    plan, decomp = explore_q.get_max();
    if decomp tasks can be added to plan then
        new_plan = plan + decomp tasks;
        add new_plan to node_collection;
        for edge in edge_collection do
            if edge.task not in new_plan and
                new_plan.cost + edge.cost below
                max_cost then
                add (new_plan, edge) to explore_q;
            end
        end
    end
end
/* find best plan in node
collection */
best_plan = null;
for plan in node_collection do
    if plan.utility above best_plan.utility then
        best_plan = plan;
    end
end
return best_plan;

```

rithm 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 $h(P, T) = P.utility + \frac{T.utility}{T.cost}$. 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.

Empirical Evaluation

To test our model, we ran simulations of our planning and execution system on three variants of the Europa Lander domain described in Figure 1. The first is the base scenario. Here each task consumes an amount of energy that matches its a priori expectation in the task network, but may be noisy, with a standard deviation of 10%. In the second variant, we bias this noise such that tasks are expected to consume 10% more energy than modeled. Finally, the third variant biases noise in the opposite direction, such that tasks are expected to consume 10% less energy. For each variant, we simulated each of the four planning/execution strategies discussed in our theoretical framework, and measured the utility achieved. In simulation, the failure probability of each task is uniform and independent. Each failure resolution mechanism is assumed to have a fixed cost and always succeed in resolving the issue. The data for each figure shows the mean utility achieved across 50 simulations of the scenario.

For our model calculations, we estimate our average utility per cost (u_{avg}) by analyzing plans generated by a precient planner. This planner has perfect execution information a priori, so plan execution exactly matches the planner's predictions. Task failure probability is assumed to be 0.1 for $P(fail) = .1$, and we assume flexible execution is able to handle 30% of such failures for $P(FE) = .03$, while replanning is able to handle an additional 60% of remaining failures for $P(replan) = 0.042$ and ground can solve the rest for $P(ground) = 0.028$.

Our model predicts the "static" strategy to perform poorly,




Figure 3: Average utility achieved in simulation of the base Europa Lander domain for 4 planning strategies, compared to theoretical model predictions.

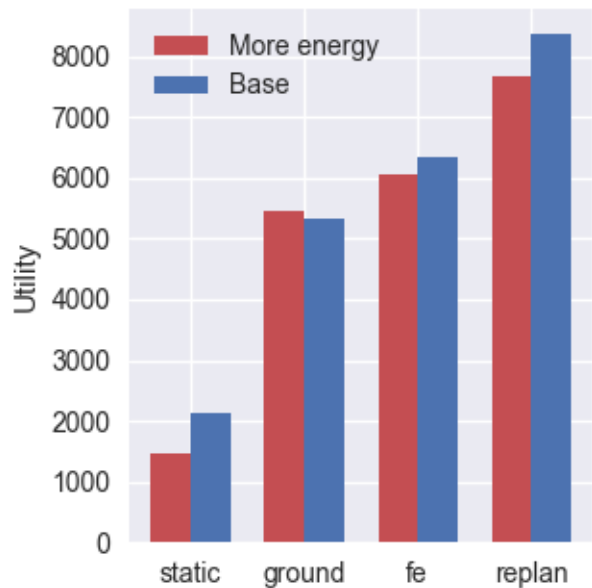


Figure 4: Average utility achieved in simulation of the Europa Lander domain where all tasks take 10% more energy than expected, compared to empirical results in the base domain.

since it has no failure resolution mechanisms and is thus likely to terminate quickly. By introducing a failure recovery mechanism, our model predicts the "ground" strategy to improve performance considerably. However, this failure re-

covery mechanism is still fairly costly. The “FE” strategy introduces flexible execution to mitigate this. As such, our model predicts a higher utility achievement, since some set of failures are now resolved by a less costly mechanism. Finally, the “replan” strategy is predicted to perform best of all the strategies. Like the “FE” strategy, it introduces another failure resolution mechanism and also introduces additional utility through plan optimization. When utility is discovered at execution time, the “replan” strategy is able to exploit that discovery, where the other strategies are not.

In Figure 3, we compare the predictions of our model to the measured utility achievement of our system in simulation. We see that the four strategies qualitatively follow the model’s predictions.

Discussion of Model Limitations and Future Model Work: While our predictive model generally matches our empirical measurements, it is limited in certain aspects. Extending the utility model to increase its accuracy would have many benefits - most importantly highlighting how different operations strategies could improve mission return.

The model uses u_{avg} as a way to estimate utility achievement based on power, smoothing performance across the entire execution into a linear model. However, in the Europa Lander domain, utility is primarily achieved only during communication events. Because the model views utility gain as purely linear, it is unable to capture the spikes in utility inherent in the domain.

In the Europa Lander domain, a trench only need to be excavated a single time, and multiple samples can be taken from a single excavation trench. This means that the first sample taken at a trench is much more costly than future samples. Because of this, if the system tends to run out of energy while attempting to sample a site for the first time, the model is likely to overestimate utility gain, since a significant portion of energy is used while no utility is gained. On the other hand, when the system can repeatedly sample from an existing trench, the model underestimates utility. This behavior is prominently seen in the ground and FE strategies in Figure 3. Both strategies spend a significant portion of their execution repeatedly sampling from an excavation site, leading to higher utility gain than expected during these portions of the plan execution.

Another issue is the simplistic plan structure presumed by our model. In general, plan dependencies can be considered a set of graphs. For the Europa Lander domain, there are not merges (except for activities shared for efficiency such as downlinks), therefore plan structures look like forests or sets of trees. In our analysis model we presume that all activities in the plan form a single linear sequence. Because a plan is a set of trees, a failure in one tree would not stop execution in another tree - reducing the failure cost. Additionally, exogenous conditions such as Earth-in-view are required for downlink so that failure costs are non linear (a 1h delay could push a downlink to the next Earth-in-view 42 hours later or conversely a delay might not delay downlink at all as the downlink was waiting for a later Earth-in-view).

For the replan strategy, we also consider the effects of utility discovery and plan optimization in replanning. To determine a value for d , the number of times that utility dis-

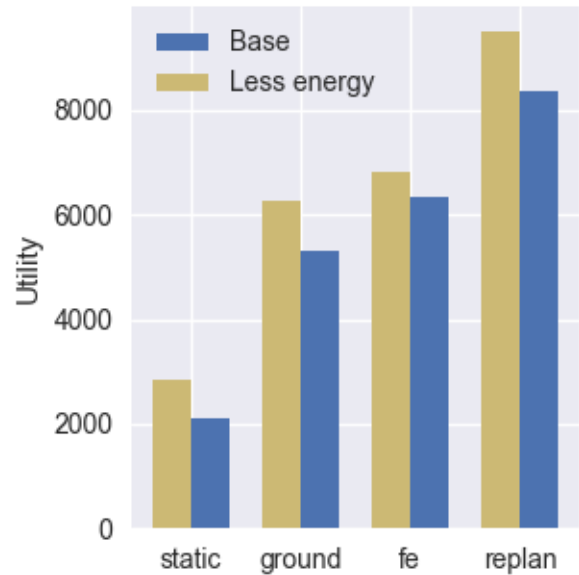


Figure 5: Average utility achieved in simulation of the Europa Lander domain where all tasks take 10% less energy than expected, compared to empirical results in the base domain.

covery can be exploited, we calculate an upper bound for this value based on the total energy available to the system. However, the system may not be able to take advantage of utility discovery this number of times, since it may run into too many task failures, or the planner may simply choose to complete other tasks. Thus, the calculations for our model tend to overestimate the value of utility discovery in the replanning strategy.

Next, we consider the effects of biased noise on the utility gain of our system. First, we examine the scenario where all tasks use 10 percent more energy on average than expected. A comparison of this scenario and the base scenario is shown in Figure 4. Naively, we might expect utility in each scenario to decrease by about 10 percent. However, because utility is achieved in spikes through the completion of fairly lengthy chains of tasks, events have an impact on utility only if they increase or decrease the probability of successfully completing a chain of tasks. In the “more energy” scenario, the ground strategy appears generally unaffected.

The replan strategy is affected more heavily, since a lower pool of energy available limits the strategy’s ability to take advantage of discovered utility. On the other hand, because it is able to replan, it can make use of lower cost actions such as Seismograph/Panorama tasks to gain utility despite lacking the energy to complete a sample.

Finally, we consider the scenario where tasks take 10 percent less energy than expected (Figure 5). Here, the ground strategy improves considerably in performance, while FE improves at a lower clip. This is consistent with what we see in the previous scenario. The ground strategy is able to

benefit significantly from the extra energy and complete an extra sample cycle, while FE is not as close to this boundary and thus is not affected as strongly.

The replan strategy also sees significant benefits from extra energy. Extra energy enables additional samples, whose benefit is amplified by the potential for utility discovery. In addition, the replan strategy is able to integrate knowledge of the additional energy during execution time as it updates state predictions with the reality on the ground. Thus, instead of settling for a Seismograph/Panorama task, as might occur in the base or high energy use scenarios, the replan strategy is more often able to process a sample.

Related Work

Decision-theoretic planning provides a formal model for reasoning about problems in which actions have stochastic outcomes or the agent has incomplete information about its environment (Iocchi et al. 2016; Saisubramanian, Zilberstein, and Shenoy 2017; Zilberstein et al. 2002). The primary objective of decision-theoretic planning is to produce plans or policies that define the potential trajectories of actions that the agent may take which maximizes its expected utility, rather than maximizing or guaranteeing goal-reachability (Boutilier, Dean, and Hanks 1999). A standard approach in decision-theoretic planning for modeling domains is to use a Markov decision process (MDP) (Bellman 1957) when the agent knows the full evaluation of every state at each timestep, or a partially observable Markov decision process (POMDP) (Spaan 2012) where this holds only for a subset of the variables that define the statespace.

However, several issues in spacecraft or rover operations complicate the use of said decision making models. First, these models traditionally do not support durative or concurrent actions, but rather assume that all actions are instantaneous and fully sequential in nature. Second, although there have been a number of approaches over the years aimed at improving the scalability of these approaches (Guestin et al. 2003; Wray, Witwicki, and Zilberstein 2017; Yoon, Fern, and Givan 2007), most algorithms that solve MDPs produce policies that account for all contingencies and provide actions for all states in the domain. This is generally impractical or impossible in spacecraft and rover operations where computational power is (often severely) limited, and more so in our problem where the battery is non-rechargeable and the domain model is expected to be modified repeatedly throughout the agent's operation.

Onboard planning and execution are of great interest to the space domain. Flexible execution of tasks is a central focus of execution engines like PLEXIL (Verma et al. 2005) and TRACE (de la Croix and Lim 2020). The Earth Observing One (EO-1) spacecraft (Chien et al. 2005), which flew for over 12 years from 2004-2017, responded to dynamic scientific events using the CASPER planner (taking 10s of minutes to replan and) with the SCL executive. The flight and ground planners (Chien et al. 2010) both used a domain specific search algorithm that enforced a strict priority model over observations for a limited model of utility.

The M2020 Perseverance rover also plans to fly an onboard planner (Rabideau and Benowitz 2017) to reduce lost

productivity from following fixed time conservative plans (Gaines and et al 2016). Like the planning approach we propose in this paper, the M2020 planning architecture also relies on rescheduling and flexible execution (Agrawal et al. 2021a), ground-based compilation (Chi et al. 2019), heuristics (Chi, Chien, and Agrawal 2020), and very limited handling of planning contingencies (Agrawal et al. 2021b). However, it uses a non-backtracking planner, which limits its ability to optimize plans and the M2020 flight software does not support utility discovery. Our work also takes a different focus, primarily examining the effects of task failure and considering integrated planning in the context of failure resolution. Finally, the Europa Lander mission concept has stronger drivers for mission autonomy than M2020 due to lack of reliable a priori model parameters, the inability to recharge the battery, and the long communications blackout time windows.

Future Work and Conclusions

Our work on FE and replanning is reactive to execution variation and failure. Likewise it is also only reactive to positive events (such as early completion or underconsumption of resources). A proactive approach would prefer plans that are resilient to negative execution outcomes or could take advantage of possible positive outcomes such as the sampling-based approach by Basich et al. (Basich et al. 2021).

In addition, in this work we focus primarily on energy as a resource. However, a number of other resources exist, and the consumption of any of these may be noisy or biased, affecting plan execution. In our application, time is a resource with complex effects (due to exogenous conditions like Earth in view for communication) and time is typically linked to energy usage (a task which runs long consumes extra energy). Data volume is another resource that can vary (science data products have variable size due to content dependent compression) and data volume translates directly to energy required to downlink.

Additionally, certain failures have effects beyond task failure. Failures in excavation or sampling damage hardware - reducing future effectiveness or even precluding activities. A true decision theoretic planner could reason about such risks and alter behavior appropriately. Other risks are due to resource uncertainty. Estimation of remaining battery energy is inaccurate; aversion to resource risk would encourage downlinking acquired science data earlier to avoid risk of unexpected early energy depletion.

We have presented a framework for analyzing the benefit of onboard autonomy for space missions. We consider alternative strategies of "go to ground", "flexible execution", and "replanning" to enable the flight system to resolve issues onboard. We describe this framework in the context of a Europa Lander mission concept and the Mexec software which uses flexible execution, HTN planning, and scheduling. We demonstrate that a predictive model of utility gain from using onboard autonomy qualitatively matches empirical results from simulation but significant further work is needed to refine and extend the model.

Acknowledgments

This work was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

- Agrawal, J.; Chi, W.; Chien, S. A.; Rabideau, G.; Gaines, D.; and Kuhn, S. 2021a. Analyzing the Effectiveness of Rescheduling and Flexible Execution Methods to Address Uncertainty in Execution Duration for a Planetary Rover. *Robotics and Autonomous Systems*, 140 (2021) 103758.
- Agrawal, J.; Chi, W.; Chien, S. A.; Rabideau, G.; Kuhn, S.; Gaines, D.; Vaquero, T.; and Bhaskaran, S. 2021b. Enabling Limited Resource-Bounded Disjunction in Scheduling. *Journal of Aerospace Information Systems*, 18:6: 322–332.
- Basich, C.; Wang, D.; Chien, S.; and Zilberstein, S. 2021. A Sampling-Based Optimization Approach to Handling Environmental Uncertainty for a Planetary Lander. In *International Workshop on Planning and Scheduling for Space (IWSS)*.
- Bellman, R. 1957. A Markovian decision process. *Journal of Mathematics and Mechanics*, 679–684.
- Boutillier, C.; Dean, T.; and Hanks, S. 1999. Decision-theoretic planning: Structural assumptions and computational leverage. *Journal of Artificial Intelligence Research (JAIR)*, 11: 1–94.
- Chi, W.; Agrawal, J.; Chien, S.; Fosse, E.; and Guduri, U. 2019. Optimizing Parameters for Uncertain Execution and Rescheduling Robustness. In *International Conference on Automated Planning and Scheduling (ICAPS 2019)*. Berkeley, California, USA.
- Chi, W.; Chien, S.; and Agrawal, J. 2020. Scheduling with Complex Consumptive Resources for a Planetary Rover. In *International Conference on Automated Planning and Scheduling (ICAPS 2020)*. Nancy, France.
- Chien, S.; Sherwood, R.; Tran, D.; Cichy, B.; Rabideau, G.; Castano, R.; Davis, A.; Mandl, D.; Frye, S.; Trout, B.; et al. 2005. Using autonomy flight software to improve science return on Earth Observing One. *Journal of Aerospace Computing, Information, and Communication*, 2(4): 196–216.
- Chien, S.; Tran, D.; Rabideau, G.; Schaffer, S.; Mandl, D.; and Frye, S. 2010. Timeline-based space operations scheduling with external constraints. In *Twentieth International Conference on Automated Planning and Scheduling*.
- de la Croix, J.-P.; and Lim, G. 2020. Event-Driven Modeling and Execution of Robotic Activities and Contingencies in the Europa Lander Mission Concept Using BPMN. In *International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS)*.
- Gaines, D.; and et al. 2016. Productivity challenges for Mars rover operations. In *Proceedings of 4th Workshop on Planning and Robotics (PlanRob)*, 115–125. London, UK.
- Goldman, R. P.; and Kuter, U. 2019. Hierarchical Task Network Planning in Common Lisp: the case of SHOP3. In *Proc 12th European Lisp Symposium*, 73–80. Zenodo.
- Guestrin, C.; Koller, D.; Parr, R.; and Venkataraman, S. 2003. Efficient solution algorithms for factored MDPs. *Journal of Artificial Intelligence Research (JAIR)*, 19: 399–468.
- Hand, K. P. 2017. *Report of the Europa Lander science definition team*. National Aeronautics and Space Administration.
- Iocchi, L.; Jeanpierre, L.; Lazaro, M. T.; and Mouaddib, A.-I. 2016. A practical framework for robust decision-theoretic planning and execution for service robots. In *International Conference on Automated Planning and Scheduling (ICAPS)*.
- Nau, D. S.; Au, T. C.; Ilghami, O.; Kuter, U.; Murdock, J. W.; Wu, D.; and Yaman, F. 2003. SHOP2: An HTN Planning System. *Journal of Artificial Intelligence Research*, 20: 379–404.
- Rabideau, G.; and Benowitz, E. 2017. Prototyping an On-board Scheduler for the Mars 2020 Rover. In *Intl Workshop on Planning and Scheduling for Space (IWSS 2017)*. Pittsburgh, PA.
- Saisubramanian, S.; Zilberstein, S.; and Shenoy, P. 2017. Optimizing electric vehicle charging through determinization. In *International Conference on Automated Planning and Scheduling (ICAPS) Workshop on Scheduling and Planning Approaches*.
- Spaan, M. T. 2012. Partially observable Markov decision processes. In *Reinforcement Learning*, 387–414. Springer.
- Verma, V.; Estlin, T.; Jónsson, A.; Pasareanu, C.; Simmons, R.; and Tso, K. 2005. Plan execution interchange language (PLEXIL) for executable plans and command sequences. In *International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS)*.
- Verma, V.; Gaines, D.; Rabideau, G.; Schaffer, S.; and Joshi, R. 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)*. Pittsburgh, PA.
- Wray, K. H.; Witwicki, S. J.; and Zilberstein, S. 2017. Online decision-making for scalable autonomous systems. In *Intl Joint Conference on Artificial Intelligence (IJCAI)*.
- Yoon, S. W.; Fern, A.; and Givan, R. 2007. FF-Replan: A baseline for probabilistic planning. In *International Conference on Automated Planning and Scheduling (ICAPS)*, volume 7, 352–359.
- Zilberstein, S.; Washington, R.; Bernstein, D. S.; and Mouaddib, A.-I. 2002. Decision-theoretic control of planetary rovers. In *Advances in Plan-Based Control of Robotic Agents*, 270–289. Springer.