Field Experiments in Mission-Level Path Execution and Re-Planning

Paul TOMPKINS, Anthony STENTZ, William WHITTAKER
The Robotics Institute, Carnegie Mellon University
5000 Forbes Avenue, Pittsburgh, PA 15213

Abstract. A new project called Life in the Atacama (LITA) motivates navigational autonomy over 100 kilometers and over days or weeks. LITA is chartered to develop technology to enable robotic astrobiology, while at the same time conducting useful research into the extremes of life on Earth. Two years from now, the team must demonstrate coordinated robotic science over several weeks. A key component is Mission-Level Path Planning, which plans routes that coordinate large-scale terrain avoidance, path timing, and battery energy management. Its goal is to yield plans that provide a framework not only for locomotion, but for other activities like solar charging and hibernation. In April 2003, we conducted the first of three field experiments in the Chilean desert. This paper documents our findings on this new level of robotic autonomy.

Introduction

The Life in the Atacama (LITA) project seeks to develop technology in support of robotic astrobiology for NASA while conducting useful Earth science in the Atacama Desert of northern Chile [13]. The Atacama is one of the driest places on Earth, and has long been known to support very little life. The project intends three seasons of rover field-testing and biological investigation, and culminates in a multi-week field demonstration in which scientists in the United States will direct a robotic search for life in Chile. To stress robot autonomy, scientists and engineers will be limited to short windows of communication, and transmission delays similar to those between Earth and Mars. Their mission will be to characterize the presence and distribution of microscopic life over more than 100 kilometers of travel in an effort to better understand the limitations of life in the Atacama ecosystem.

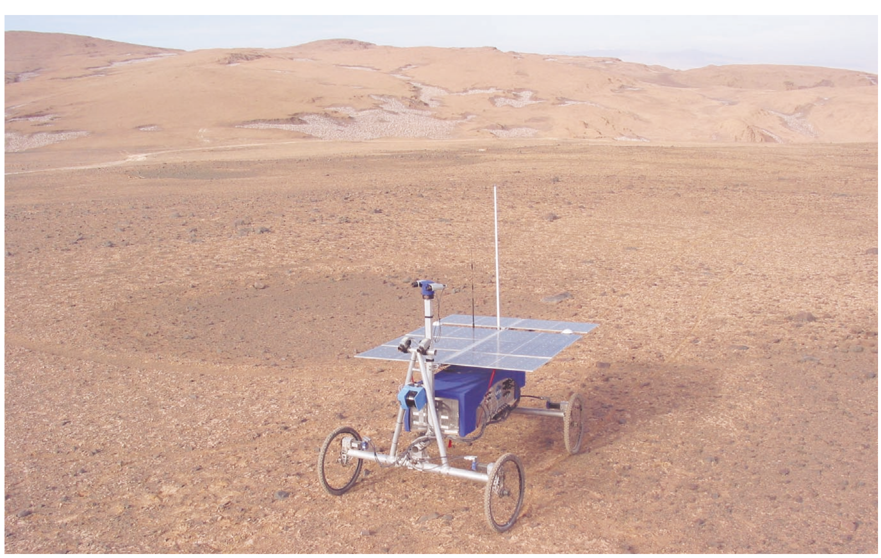

Figure 1: Hyperion in the Salar Grande, Atacama Desert, April 2003.

The goals of the LITA field investigation dictate traverses between distant science survey sites, opportunistic science, and extended 24-hour autonomy, including re-configuration for hibernation during the night. The mission profile is an ideal application of Mission-Level Planning - navigational autonomy that considers the long-range effects of actions in terms of route, timing and resources. In contrast to local navigation that considers spans of meters and minutes at high resolution, mission-level path planning considers hundreds of meters and hours at lower resolution. Rather than plan for locomotion exclusively, a mission-level path planner considers how certain other events, such as stationary solar charging or hibernation, affect a day's mobility activities. In bridging the gap between path planning and classical planning and scheduling, a mission-level navigation planner could provide a coarse framework, based on an integrated approach, from which to derive more detailed plans.

In April 2003, building upon technology developed in the Sun-Synchronous Navigation project [12], the first LITA field experiment used the Hyperion robot (Figure 1) to test a number of technologies including mission-level path planning. Over eight days, TEMPEST provided plans and re-plans for Hyperion that avoided hazardous terrain, considered the solar energy available to the robot, and satisfied operational constraints.

We begin by briefly describing our planning approach, then examine a portion of the planning and re-planning results from Atacama field experiments. We compare our work to other relevant robot field work next, and conclude with a critique of this year's accomplishments, and suggestions of what needs to be done in upcoming years.

1. Approach

1.1. Hyperion and Autonomy Architecture

Hyperion is a solar-powered robot originally designed for sun-synchronous navigation in polar latitudes [12]. The Atacama Desert differs from the Arctic in a number of respects, most notably in latitude, which affects sun elevation angles, diurnal lighting and the navigation strategy. For the Atacama, Hyperion's solar array was oriented horizontally to best collect energy from the overhead sun. Even so, it's batteries will last only two hours while driving without the sun's input, so it is critically dependent on its environment.

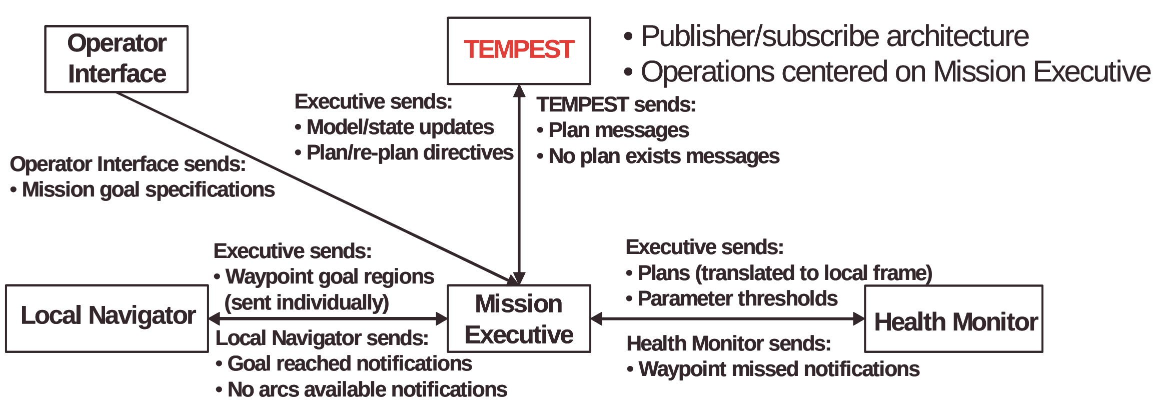

Figure 2: The Mission-Level Planning and Execution modules for Hyperion.

A principal task in the first year was to enable autonomous driving of over 1 kilometer. The Hyperion autonomy software comprises a Health Monitor (HM) in charge of responding to abnormal state conditions while operating, a Navigator that uses stereo camera data to locate and avoid hazardous local terrain while seeking a goal, and TEMPEST. A rudimentary mission executive (ME) coordinates mission-related data passing between these modules, and receives and distributes traverse commands from the Operator Interface (OI). Figure 2 illustrates the basic set of inter-module communications relating to mission-level path planning and execution. Designed as a placeholder for future, more sophisticated executive modules, the ME adopted a simple, hierarchical state machine architecture. It enacted a collection of monitors to keep track of system state while planning or executing plans: the Plan Request Monitor, Plan Execution Monitor, Drive Monitor, and Charge Monitor.

Upon receiving a traverse command from the OI, the ME determines the current robot position and time state, and then requests a plan from the initial state to the traverse goal state. TEMPEST finds and returns an optimal plan to the ME. For each action in the plan, the ME triggers one of two Action Monitors.

The Drive Monitor computes parameters for a 10 meter by 30 meter goal region surrounding the next position waypoint (see Figure 3), and sends them to the Navigator for execution. Using the goal region as its global goal, the Navigator pursues this region while avoiding obstacles it detects. Once Hyperion is within the region, the Navigator signals its arrival, terminating the Action Monitor. The Charge Monitor stops the rover and then waits for the target Charge duration before terminating.

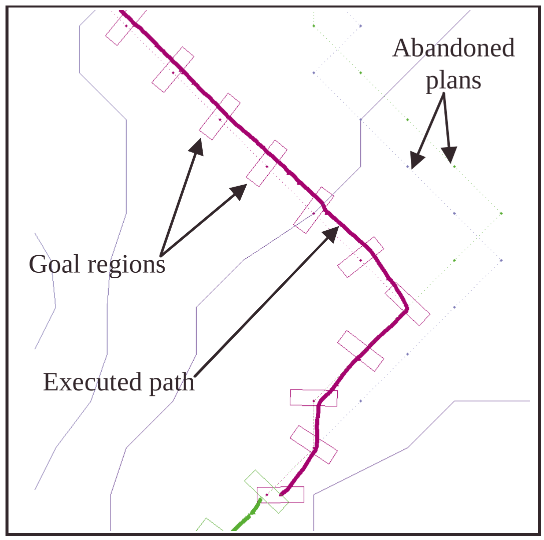

Figure 3: Plans and executed paths. Goal regions gave the local navigator flexibility in selecting the specific path between waypoints.

During plan execution, at the scheduled arrival time for each plan waypoint, the HM performs a one-time check to confirm the robot is on time. If the rover is more than a fixed distance from the waypoint at the scheduled arrival time, the HM requests a re-plan. The ME terminates the execution of the current plan, and TEMPEST uses the current rover state and previous goal to find a new plan.

1.2. TEMPEST and Incremental Search Engine

TEMPEST is a planner that considers large-scale terrain, time-varying sunlight, rover power and operational constraints in creating paths. In the context of LITA, it resides as the upper layer in two-layer navigational hierarchy. Coarse models of terrain and rover operations enable TEMPEST to project activities hours into the future to gauge the impact of hills, valleys, sunlight, shadows and finite battery capacity. TEMPEST produces plans that can guide local planners with respect to route, event timing and minimum recommended battery energy. At its core, the planner relies on an algorithm called Incremental Search Engine (ISE), which enables globally-constrained, high-dimensional path search [8][10]. More importantly, ISE provides a means of very efficient re-planning as models of the world and rover evolve. In Sun-Synchronous Navigation experiments [9], TEMPEST was run offline and in a single-shot mode. Recent improvements in TEMPEST performance, combined with the incorporation of ISE plan extension and re-planning, have yielded an effective online planner and re-planner.

As a first step toward more comprehensive re-planning, TEMPEST calls on simple plan extension to update plans to a new start state. It is essential that a planner be able to adjust a plan in response to deviations from an earlier plan. Perhaps the rover falls behind schedule, or is forced to divert from a previously planned route to avoid rocky terrain that does not appear in the terrain model. ISE handles these cases very easily by simply extending its existing search graph to the new state, and re-applying its feasibility criteria to candidate solutions arriving there. This is simpler than full re-planning because there is no need to repair the search graph - the underlying assumptions from the previous search remain valid.

Despite the simplicity of this new addition, it brings tremendous utility to TEMPEST. Where during the Sun-Synchronous Navigation experiment plans were generated once for a 24-hour traverse, and relied on perfect execution to remain valid, TEMPEST can now repair plans as the execution evolves.

TEMPEST plans are sequences of Drive, Charge and Hibernation actions predicted to follow a specifically scheduled route while maintaining sufficient battery charge and satisfying global constraints. The plan follows “waypoints” through a 4-D state space of x, y, time and battery energy. We refer readers to the companion paper, in these conference proceedings, for further details on TEMPEST [10].

2. Field Experiment

Over April 17 through 20 and April 24 through 26, TEMPEST generated 27 plans and 83 replans. Our goal was to evaluate TEMPEST in terms of resulting path length, terrain avoidance, energy and operational utility.

2.1. Path Length

The ISE grid representation of position state interferes with the ability to produce shortest-distance paths. The grid enforces motion on an eight-connected graph connecting each cell to its nearest neighbors. The minimum eight-connected plan distance between two points depends on the ratio of their relative distance in the x-coordinate *(Δx)* and the y-coordinate *(Δy)*. When the start and goal lie along the horizontal ($\Delta x/\Delta y$=infinity), vertical ($\Delta x/\Delta y$=0) or principle diagonal axes of the graph ($\Delta x/\Delta y$=1 or $\Delta x/\Delta y$=-1), the minimum eight-connected plan distance is equal to the Euclidean distance between the points. Since the eight-connected path cannot assume arbitrary headings, other ratios of Δx-to-Δy produce minimum path lengths that exceed the Euclidean distance between the points.

In Figure 4 we examine plans generated during the field experiment to determine the degree to which eight-connectedness was dominant in extending path length beyond the Euclidean distance. The horizontal axis spans the range of the absolute value of $\Delta x/\Delta y$, representing an East-West heading on the left, a Northeast-Southwest or Northwest-Southeast heading at the center, and a North-South heading on the right. The vertical axis spans a range of the plan distance divided by the Euclidean or map distance. The blue curve shows the minimum plan-to-map distance ratio for the range of ratios of Δx and Δy. The markers indicate the relationship between grid distance ratio and plan-to-map ratio for all the plans generated on April 20 through April 26. One might first observe that all the plans fall to the right of center, confirming that routes traveled principally in a North-South direction. More interestingly, though several plans fall very close to the minimum curve, most paths are much longer. This indicates that the eight-connected effect was not the principal contributor to path length extension for most of the plans. For the paths not on the

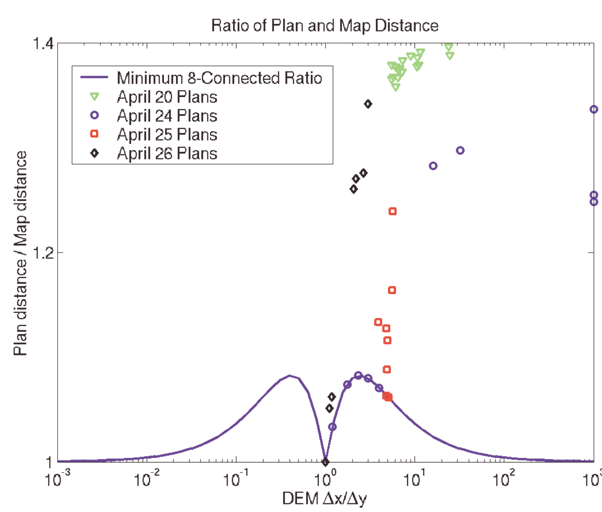

Figure 4: The eight-connected path artifact was not often dominant in determining path length.

minimum eight-connected curve, obstacle avoidance, energy cost minimization and constraint satisfaction contributed to path length extension, often significantly.

2.2. Large-Scale Terrain Avoidance

TEMPEST demonstrated large-scale hazard avoidance on several occasions. The planning for April 25 shows subtlety. Figure 4a shows the sequence of plans and executed paths for the day. At first glance, the initial northeast heading taken by the plans is mysterious. Why did the planner force this detour rather than a more direct route to the goal? The answer appears to lie in slope avoidance. By plotting the same path over a contour map of the magnitude-of-gradient (slope) field (see Figure 6b), we observe that the path avoids steeper slopes to its left, and then turns toward the goal at a break in this higher slope region.

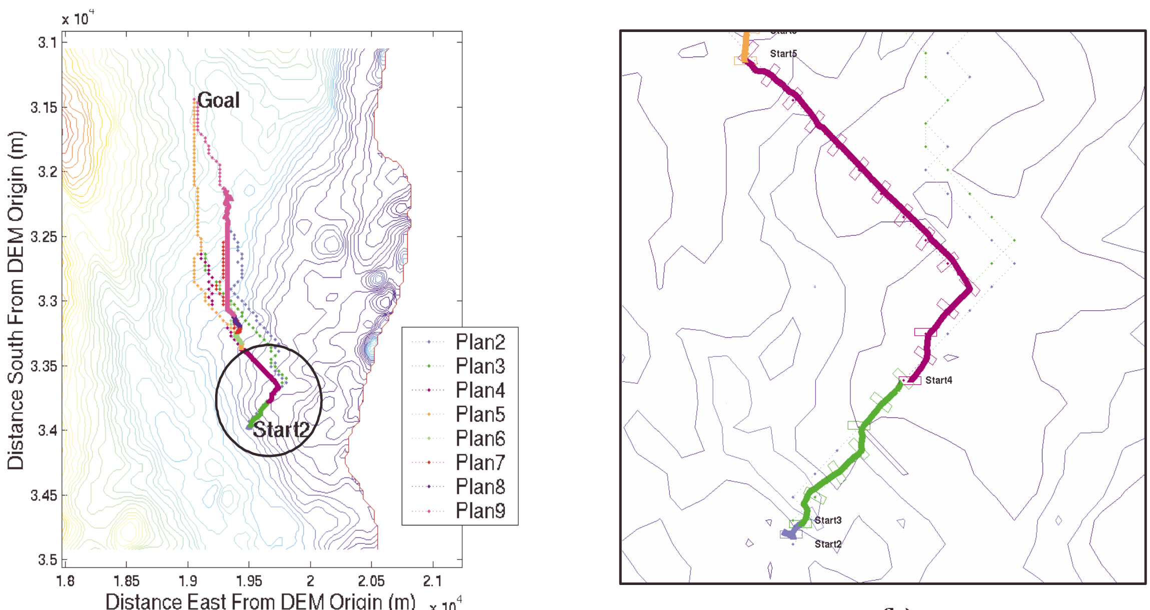

(a) (b)

Figure 5: Plan and re-plan routes from April 25 on an elevation contour map, and a close-up with contours of constant slope. The initial plans seem to have located a break in steeper slopes.

The most interesting, yet greatest failure in terrain avoidance occurred on April 18 (see Figure 6). In a plan to travel to the southern end of the area of operations, TEMPEST dictated a traverse near the rim of the large fault running in a primarily North-South direction to the West of base camp. According to observers near the robot on this day, waypoint goals dictated travel down precariously steep slopes on the West side of the fault ridge. This motion prompted the team to abort autonomous travel at this point. To this date, it is unclear why waypoints forced the rover onto dangerous terrain.

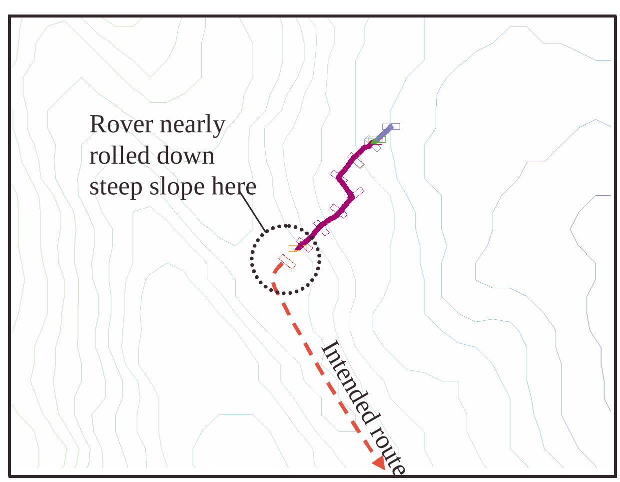

Figure 6: TEMPEST placed Hyperion very close to a hazardous slope on April 18. Map registration errors or position uncertainty may have been to blame.

2.3. Energy Efficiency

In contrast to previous planning experiments in the Arctic [9], evaluating the plans from the Atacama field experiment proves to be difficult. For Atacama experiments, Hyperion's solar array was horizontal. This removes the strong coupling between the direction of travel and solar power in Hyperion's Arctic configuration, allowing the rover to fall behind schedule with little penalty. Furthermore, the solar flux in the Atacama was sufficiently high in April, during daylight, to sustain the highest-power operations indefinitely. Shadows only occurred very near sunset, so only intersected paths when operations were coming to a close. The planning models verify this - TEMPEST executed plans never included Charge or Hibernation actions. Finally, due to an undiscovered software bug, the telemetry logs did not log TEMPEST plan messages. Alternate records of plans, used to reconstruct plans from April 20 and later, did not include the battery energy variable of the plans. This prevents determining when and where planning predicted energy-rich and energy-poor conditions.

2.4. Plan Monitoring and Re-Planning

A primary goal of the field experiment was to test re-planning in the context of rover operations and plan stability. As mentioned earlier, TEMPEST called upon the simplest form of ISE re-planning - updating a plan in response to a new initial state. The Health Monitor provided simple plan execution monitoring, and was the sole trigger of re-planning.

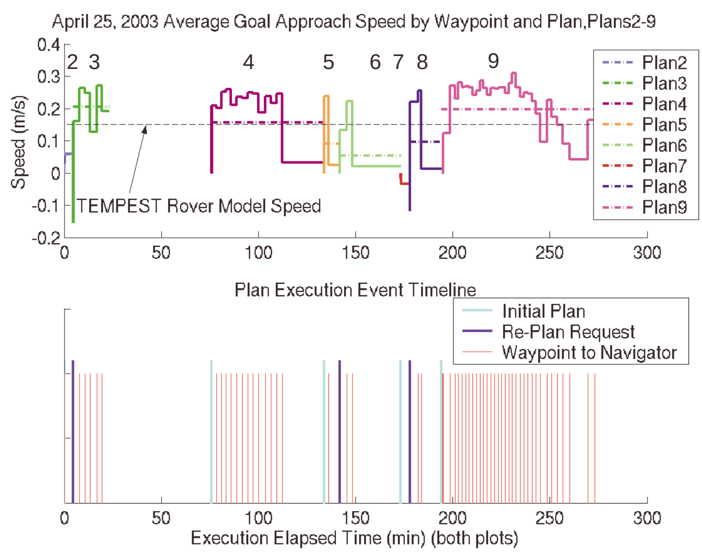

Figure 7: Rover Average Speed vs. Re-Plan Frequency. Operational delays often caused the HM to trigger re-planning. The solid lines in the upper plot show average rover speed over a Drive action. The dashed lines are the average rover speed over the particular plan or re-plan execution. Speeds below the TEMPEST rover model speed caused re-plans. Blank regions in the upper plot signify a suspension of operations.

As state estimation was quite accurate, the major cause of re-plan requests was deviation of average rover speed from the rover model, shown in Figure 7 for plans executed on April 25. The figure illustrates the connection between rover speed and re-planning triggers. Figure 7a plots speed as a function of time. The TEMPEST rover model speed is the constant dashed black line. The solid traces show the average rover speed for each executed Drive action. The dash-dot traces show the average rover speed over the entire plan execution. Meanwhile, Figure 7b shows the timing of important plan execution events. TEMPEST was terminated several times during the day after long operational delays (after Plan 3, Plan 4, Plan 6 and Plan 8). However, it is clear from Plan 2, Plan 5 and Plan 7 that re-plans correlate well with periods of slow average driving speeds. The figure also underlines a logic error in the Health Monitor that overlooked faster-than-expected rover speed for re-planning. In no case does faster-than-predicted rover speed trigger a re-plan (see Plan 9).

2.5. Plan Stability

When re-planning, it is important to analyze how successive plans compare with earlier plans. The degree to which routes are stable determines whether scientists can rely upon visiting specific interesting areas en route to the final goal. The answers might also be useful for software designers interested in integrating TEMPEST with other planning or scheduling modules. For example, if re-plans cause a total re-specification of a timeline, it may only be practical to create a detailed command schedule for the first action in the plan, while stable re-planning may permit a longer projection to remain valid.

Figure 8a illustrates plan route stability for the field experiment. For a given re-planning instance, the horizontal axis shows the fraction of the re-plan waypoints that are identical to the initial plan (F_i). The vertical axis shows the fraction of the re-plan waypoints that are identical to those from the previous plan (F_p). Each trace represents a series of re-plans to a single goal. The markers on the traces correspond to results from a specific re-planning instance. The traces begin at the enlarged markers, the first re-plan, and go in chronological order. It follows that all traces begin on the line $F_i = F_p$, since for the first re-plan, the previous plan is also the initial plan.

We observe that for all but one trace, the endpoint falls generally left and above the starting point. One can infer that for these cases, re-plans are initially unstable but grow gradually more stable as plan execution progresses. This seems to make intuitive sense. With the greater freedom that comes with a large distance between start and goal, ISE finds a number of plans of similar cost but with differing routes. Subtle changes in initial conditions may cause substantial route variations. However, as the distance to the goal shrinks, the freedom is reduced, causing greater stability.

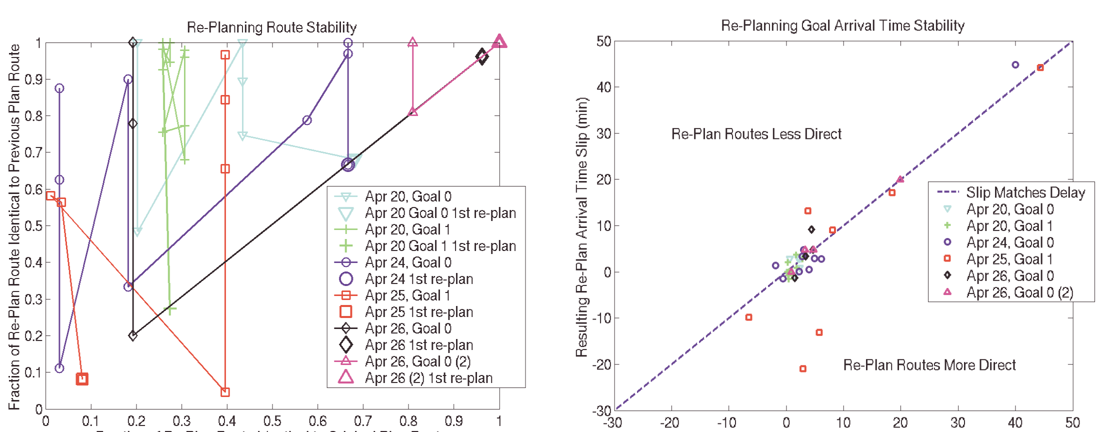

(a) (b)

Figure 8: Route Stability and Arrival Time Stability

The exception is the plan sequence from April 25, whose first re-plan shares fewer than 10% of the initial route's waypoints. Successive re-plans deviate even more from the initial plan at first, but then return to match about 40% of the remaining plan. Figure 3a may help clarify what is happening in this case. The plans seem to alternate between two general routes over the last 1/2 of the traverse. Plan 2 (the initial plan shown) takes the right fork, Plan 3 (the first re-plan) the left. In the first half of the route, Plan 2 and Plan 3 are almost entirely distinct, but run very close to each other. In later planning instances, the plans settle on a variation of the right fork, increasing the fraction of the plan that is identical to Plan 2.

We also investigate plan arrival time stability. We wanted to determine whether projected goal arrival time slips were explained well by the delays induced by lower-than-anticipated rover speeds. In Figure 8b, we plot arrival time slip against operational delays. Each marker corresponds to a different plan instance. The dashed line falls where schedule slip would exactly match rover delays. Re-plans falling above the dashed line are less direct then their predecessors, while re-plans below the line are more direct. Aside from a few outliers, the data seems to suggest a strong correlation between operational delays and schedule slips.

3. Related Work

Several teams have conducted long-range rover navigation field experiments in planetary analog terrain. Most research has focused on local navigation to provide safe, goal-directed travel over tens or hundreds of meters [2][3][4][5][7][11]. We view all of these developments as enabling for mission-level path planning - without local navigation, our technique could not work. However, they have substantially different research goals than our own. Path planning and activity scheduling are almost universally functionally separated from each other. However, path planners and classical AI planner-scheduler software is becoming more tightly coupled, for example as demonstrated on a rover in the work of Estlin et al. [1].

Relating to energy-cognizant planning, Shillcut [6] analyzed the motion of the sun to select amongst coverage patterns for meteorite search in the Antarctic, but did not incorporate these analyses into path planning. TEMPEST planned multi-kilometer cyclic routes off-board for open-loop execution by Hyperion [9]. However, it did not re-plan, and due to slow performance, planning was limited to route timing, and path selection for 400 meter intervals between pre-selected waypoints.

4. Conclusions

We believe the greatest contribution of this work to be the operational proof-of-concept for TEMPEST. The two-level navigation hierarchy restricts TEMPEST to work at a coarse, abstract level, at large scale, without knowledge of local terrain features. Goal regions allowed the Navigator to follow the plan, yet provided the flexibility to allow adaptation to local terrain. TEMPEST naturally and efficiently re-planned routes during execution, re-synchronizing in response to delays. Now that re-planning has proven useful, augmenting world or rover models using sensor measurements might yield some very interesting results. Specifically, rover plan execution could benefit immediately by incorporating sensor measurements of terrain into TEMPEST models. Many of the terrain avoidance failures might have been sidestepped with improved ability to perceive, identify and model mid-scale features such as extended slopes and rocky outcrops. Inserted into flexible terrain maps, these measurements could fold smoothly into the TEMPEST re-planning model.

Unfortunately, there is little data to evaluate rover energy modeling. The Atacama's high solar flux and the insensitivity of Hyperion solar power to orientation might point to de-emphasizing energy optimization for future years. However, the next field experiment will entail 24-hour autonomous operations. Dawn, dusk and nighttime may force extended battery charging and low-power hibernation. From the perspective of planning, it will be important to add a battery charge estimator to Hyperion, and to develop a strategy that considers this estimate when executing plans.

Evidence suggests the rover should be more robust to sources of model error and uncertainty. Rover model speed errors prevent accurate projections of schedule, which in turn reduces the accuracy of energy projections. However, it is not clear TEMPEST will be able to adequately compensate for the full range of possible operational delays seen in the first year. More realistically, TEMPEST might be able to plan for the speed uncertainty caused by unknown local obstacle density. Further, machine learning might be used to adapt the TEMPEST rover model to changing conditions.

Avoiding terrain hazards will be critical in future years. Map registration errors may exist in future elevation maps, and position estimation may become a substantial problem with multi-day autonomy. Again, TEMPEST may be able to address these problems, or might rely more on terrain sensing to avoid mishaps.

Currently, science activities are totally decoupled from TEMPEST planning. Science activities were interleaved manually with plans, but neither TEMPEST nor the mission executive could predict their impact on schedule or energy resources ahead of time. Given the project's emphasis on scientific investigation, it will be important to consider science activities both in planning and execution.

References

[1] T. Estlin, R. Castano, R. Anderson, D. Gaines, F. Fisher, M. Judd, “Learning and Planning for Mars Rover Science,” *Workshop on Issues in Designing Physical Agents for Dynamic Real-Time Environments: World Modeling, Planning, Learning, and Communicating, International Joint Conference on Artificial Intelligence (IJCAI 2003)*, Acapulco, Mexico, August 2003.

[2] S. Lacroix, A. Mallet, D. Bonnafous, G. Bouzil, S. Fleury, M. Herrb, R. Chatila, “Autonomous Rover Navigation on Unknown Terrains: Functions and Integration”, *International Journal of Robotics Research*, 21(10-11), pages 917-942, 2002.

[3] S. Laubach, J. Burdick, “RoverBug: An Autonomous Path-Planner for Planetary Microrovers,” *Sixth International Symposium on Experimental Robotics (ISER 99)*, Sydney, Australia, March 1999.

[4] S. Moorehead, R. Simmons, D. Apostolopoulos, W. Whittaker, “Autonomous Navigation Field Results of a Planetary Analog Robot in Antarctica,” *Proceedings of the 5th International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS 99)*, June, 1999.

[5] H. Seraji, A. Howard, “Behavior-Based Navigation on Challenging Terrain: A Fuzzy Logic Approach,” *IEEE Transactions on Robotics and Automation*, 18(3), pp. 308-321, June 2002.

[6] K. Shillcut, “Solar Based Navigation for Robotic Explorers”, Ph.D. thesis, Carnegie Mellon University Robotics Institute Technical Report CMU-RI TR-00-25, October 2000.

[7] A. Stentz, M. Hebert, “A Complete Navigation System for Goal Acquisition in Unknown Environments”, *Autonomous Robots*, Vol. 2, No. 2, August 1995.

[8] A. Stentz, “Optimal Incremental Search for High-Dimensional, Constrained Path Finding Problems,” Carnegie Mellon University Robotics Institute Technical Report, to be released.

[9] P. Tompkins, A. Stentz, W. Whittaker, “Mission Planning for the Sun-Synchronous Navigation Field Experiment,” *Proceedings of the 2002 IEEE International Conference on Robotics and Automation (ICRA ‘02)*, Washington D.C., May 2002.

[10] P. Tompkins, A. Stentz, W. Whittaker, “Mission-Level Planning for Rover Exploration,” *Proceedings of the 8th Conference on Intelligent Autonomous Systems (IAS 04)*, Amsterdam, Netherlands, March 2004.

[11] E. Tunstel, T. Huntsberger, H. Aghazarian, P. Backes, E. Baumgartner, Y. Cheng, M. Garrett, B. Kennedy, C. Leger, L. Magnone, J. Norris, M. Powell, A. Trebi-Ollennu, P. Schenker, “FIDO Rover Field Trials as Rehearsal for the NASA 2003 Mars Exploration Rovers Mission,” *9th International Symposium on Robotics & Applications, 5th World Automation Congress,* Orlando, FL, 2002.

[12] D. Wettergreen, B. Dias, B. Shamah, J. Teza, P. Tompkins, C. Urmson, M. Wagner, W. Whittaker, “Experiments in Sun-Synchronous Navigation,” *Proceedings of the 2002 IEEE International Conference on Robotics and Automation (ICRA ‘02)*, Washington D.C., May 2002.

[13] D. Wettergreen, N. Cabrol, F. Calderon, M. Deans, D. Jonak, A. Luders, F. Shaw, T. Smith, J. Teza, P. Tompkins, C. Urmson, V. Verma, A. Waggoner, M. Wagner, “Life in the Atacama: Field Season 2003 Experiment Plans and Technical Results,” Carnegie Mellon University Robotics Institute Technical Report, CMU-RI-TR-03-50, October, 2003.