SOFIA's Challenge: Automated Scheduling of Airborne Astronomy Observations

Jeremy Frank
NASA Ames Research Center
Mail Stop N269-3
Moffett Field, CA 94035-1000
frank@email.arc.nasa.gov

Abstract

The Stratospheric Observatory for Infrared Astronomy (SOFIA) will require scheduling flights in support of observations proposed by many different investigators. Automation will be crucial to enable efficient and effective scheduling of SOFIA flights. We have designed an Automated Flight Planner (AFP) that accepts as input a set of requested observations, designated flight days, weather predictions and fuel limitations, and searches automatically for high-quality flight plans that satisfy all relevant aircraft and astronomer specified constraints. The AFP can generate one week's worth of flights in tens of minutes of computation time, a feat beyond the capabilities of human flight planners. The rate at which the AFP can generate flights enables a small staff to assess and analyze complex tradeoffs between fuel consumption, estimated science quality and the percentage of scheduled observations.

1. Introduction

The Stratospheric Observatory for Infrared Astronomy (SOFIA) is NASA's next generation airborne astronomical observatory. The facility consists of a 747-SP modified to accommodate a 2.5 meter telescope. SOFIA is expected to fly an average of 140 science flights per year over its 20 year lifetime. Depending on the nature of the instrument used during flight, 5-15 observations per flight are expected. The SOFIA telescope is mounted aft of the wings on the port side of the aircraft and is articulated through a range of 20 to 60 degrees of elevation. The telescope has minimal lateral flexibility; thus, the aircraft must turn constantly to maintain the telescope's focus on an object during observations. A significant problem in SOFIA operations is that of scheduling flights in support of observations. Investigators are expected to propose small numbers of observations, and many observations must be grouped together to make up single flights.

Flight planning for the previous generation airborne observatory, the Kuiper Airborne Observatory (KAO), was done by hand; planners had to choose takeoff time, observations to perform, and decide on setup-actions (called "dead-legs") to position the aircraft prior to observing. This task frequently required between 6-8 man hours to plan a single flight. The scope of the flight planning problem with the anticipated flight rate for SOFIA makes the manual approach for flight planning daunting. However, automating the planning of flights contains its own challenges. In this paper, we will summarize the technical challenges that have been met in order to build an automated flight planner, and discuss future challenges.

2. SOFIA's Operations Challenge

SOFIA's predecessor, the Kuiper Airborne Observatory (KAO), was a Principle Investigator (PI) driven facility. Scientists who developed instruments had primary responsibility for developing flights in support of their own science. PIs had considerable latitude to choose individual targets, observation duration, and groups of observations on flight days, even the time of year during which to observe. However, this also forced PIs to have a flight planner on their staff.

A variety of tools were developed in order to support this mode of operations. One of these, KNAV, was a command-line driven tool that accepted as input a specification of a flight leg, and calculated the ground track of the aircraft and the object elevation during the leg. The human flight planner had to choose the takeoff time, observations to perform, duration, and decide on "dead-legs" to position the aircraft prior to observing. The human planner also had to check the details of each leg for elevation limit violations, Special Use Airspace (SUA) incursions, and estimate fuel consumption. Former KAO flight planners indicated that one flight required 6-8 hours to plan.

By contrast, roughly half of SOFIA's flights will be driven by General Investigators (GIs), who will propose small numbers of observations that must be assembled into flights by SOFIA operations staff. Flight planners will receive lists of proposals to perform observations with one instrument. Each proposal will be ranked in importance by a Time Allocations Committee (TAC), and an instrument schedule will be developed using these proposals. An instrument will stay on the aircraft for a period of weeks; instrument are changed on weekends. The expected flight rate is 2-3 flights a week. Flight planners will then build flights for each of these instrument blocks. The set of targets and the (cumulative) durations of the observations will be fully specified, providing less latitude to flight planners. Furthermore, it is expected that there will be more observations than can be scheduled, requiring tradeoffs to optimize flights. Due to the expected number of flights and complexity of the problem, scheduling flights by hand would require a large staff, with the attendant expense of training and overhead costs.

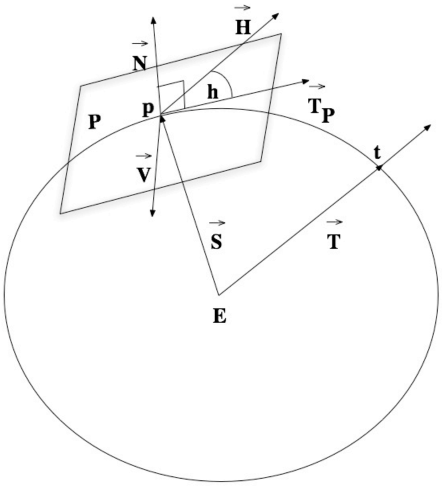

3. The Automation Challenge: Constraints

Figure 1. Solving for target elevation and ground track.

Automation will be critical in order to effectively plan flights with a small staff. Automation will not replace human flight planners, but will provide human planners with tools they can use to plan flights more efficiently, and more effectively, with smaller staff. This approach has been effective in scheduling astronomical observations for the Hubble Space Telescope [6], Earth Observing Satellites [8] and Mars Exploration Rovers [1]. However, successful automation of flight planning introduces new challenges not addressed by previously designed automated planning systems.

SOFIA consists of a telescope carried aboard a Boeing 747-SP aircraft. SOFIA can view objects between 20 and 60 degrees of local elevation (from the plane of flight). The elevation of an object depends on the object's coordinates, the aircraft's position and the time, and the combination of starting position and starting time of an observation limits the amount of time an observation can be seen continuously. Checking this constraint requires computing the aircraft's ground track throughout the course of the observation. Figure 1 shows the interaction between the object's coordinates, the aircraft's position, the time, and the telescope elevation. The Earth is shown as an oblate spheroid **E**. Let **p** be the aircraft's current position, (latitude γ and longitude L) and θ be the (Sidereal) time that the aircraft is at **p**. Let **S** be the vector from the center of **E** to **p**. Let **T** be the vector to an astronomical object **o** at time θ and **P** be the plane tangent to **E** at **p**. Let $\mathbf{T_P}$ be the projection of **T** onto **P**; this is the *object azimuth* at **p**. Let **V** be the desired heading of the aircraft. The observatory must track the object inducing **T** subject to the constraint that the angle between **V** and $\mathbf{T_P}$ is 270 degrees, because the telescope points out the left-hand side of the aircraft. Let **H** be the elevation vector with respect to **P**. Most targets are sufficiently far from Earth that we can assume $\mathbf{H} = \mathbf{T}+\mathbf{S}$. Due to the elevation limits imposed by the physical design of the telescope cavity, the angle **h** between **H** and $\mathbf{T_P}$ must be between 20 and 60 degrees throughout an observation.

Solving for the ground track is necessary to compute **h** over the entire duration of the observation and check the elevation constraints. The vector **T** traces a circle follows a latitude line around the Earth in 24 (Sidereal) hours; the latitude is the same as the observation declination δ. By contrast, the aircraft will be flying at Mach .84, which translates to roughly constant airspeed (depending on the outside air temperature). The aircraft ground speed will vary depending on the airspeed and local wind conditions. Thus, the direction of **V** is changing constantly, and its rate of change depends on both the magnitude of **V** (the ground speed) and the rate at which **T** is changing. Finding the ground track requires solving a differential equation induced by the coupled motion of **T** and **V**. The function describing the change in elevation **h** requires the ground track as input, and also depends on the rate of change of **T**.

Unlike ground-based and space-based telescopes, observing time for SOFIA is limited by flight time, which in turn is limited by the initial fuel load and the rate of fuel consumption. The fuel consumption of each engine depends on the aircraft weight, mach number, outside air temperature, initial altitude and final altitude. Since aircraft weight changes as fuel consumed, calculating fuel consumption also requires solving a differential equation. Predicted or forecast wind and temperature are interpolated and used to calculate the fuel consumption.

Infrared signal is attenuated by atmospheric water vapor. Much of the atmospheric water vapor is below the stratosphere; the goal of SOFIA is to fly above atmospheric water vapor to enable superior infrared observing compared to ground-based telescopes. Observations may include constraints on line-of-sight water vapor (LOS WV). Atmospheric water vapor varies geographically, and is generally lower over the poles; it also diminishes with increased altitude. The altitude achievable by the aircraft depends on outside air temperature and aircraft weight; the heavier the aircraft, the lower the operational ceiling of the aircraft due to the required lift. It is also possible to reduce LOS WV by observing targets at high telescope elevation; this can reduce LOS WV by a factor of 2-3.

Atmospheric conditions in the upper atmosphere during an observation are uncertain at observing time. Furthermore, the targets have uncertain properties that influence the signal received by the instrument. To account for this, targets with well known properties are used to make *calibration* measurements. Calibrations must be done periodically throughout a flight in order to ensure that, for any observation, data from a well characterized target was collected under similar conditions. This introduces new observation requests with *ordering constraints*. For instance, an astronomer may request a calibration observation no more than 30 minutes before or after a science observation.

Finally, Special Use Airspace (SUAs) constrain the ground track of the aircraft. SUAs encompass a geographic region, have a ceiling (and sometimes a floor). SUAs may also be inactive on some days. SUAs constrain the ground track of the aircraft, and thus implicitly constrain flight plans.

4. The Automation Challenge: Search

June

Name	RA (J2000)	Dec	Total time hours	Rank
Alpha Lyr	18:36:56.34	+38:47:01.3	1.50	5
Alpha Lyr	18:36:56.34	+38:47:01.3	1.50	5
Fomalhaut	22:57:39.05	-29:37:20.1	1.50	5
Fomalhaut	22:57:39.05	-29:37:20.1	1.50	5
HD 141569	15:47:20.2	-03:46:12	1.00	5
Sag A West, Arches, filaments	17:45:50.5	-28:49:28	1.74	4
Sag A West, Arches, filaments	17:45:50.5	-28:49:28	1.74	4
Sag A West, Arches, filaments	17:45:50.5	-28:49:28	1.74	4
Sag A West, Arches, filaments	17:45:50.5	-28:49:28	1.74	4
Sag A West, Arches, filaments	17:45:50.5	-28:49:28	1.74	4
"Pistol" star, Sickle region	17:46:15.3	-28:50:04	1.35	4
M83	13:37:00.78	-29:51:58.6	1.60	3
M83	13:37:00.78	-29:51:58.6	1.61	3
NGC 6946	20:34:17.0	60:08:58.0	1.48	3
NGC 6946	20:34:17.0	60:08:58.0	1.48	3
NGC 6946	20:34:17.0	60:08:58.0	1.50	3
M16	18:18:51.5	-13:49:30	1.07	2
M16	18:18:51.5	-13:49:30	1.08	2
M17	18:20:26	-16:10:36	0.70	2
W49 A	19:10:17	09:06:00	1.11	2
W51 IRS1/IRS2	19:23:41	14:31:30	0.26	2
Rho Oph	16:26	-24:20	1.02	1
Rho Oph	16:26	-24:20	1.03	1

Figure 2. A sample flight planning problem. The list of target coordinates, requested observation duration, and TAC rank are shown. These targets are to be observed over five days in June.

A flight planning problem is posed as follows: a flight planner is presented with a list of a tens to hundreds of observation requests, and a period of weeks in which to schedule the observations. Flight days are designated in advance, since most observations are performed with a particular instrument, and the instrument schedule is planned by quarters or semesters. Observations are grouped into proposals, each of which is given a TAC rank from 1-5; Rank 5 projects are most important. The task for the flight planner is to choose the observations to schedule, group them on flight days, and construct a series of flight plans for those observations. These flight plans must satisfy all of the constraints described above. Generally all rank 5 objects are expected to be scheduled; remaining time is allocated to rank 4 observations, and so on. Observations may either have constraints on permitted LOS WV achieved in a schedule, or there may simply be a desire to minimize LOS WV. A sample problem is shown in Figure 2.

It is expected that there will be too many objects to schedule, so a flight planner must search for a *subset* of observations to schedule. Furthermore, the *order* in which observations are scheduled in a flight will have an impact on both the number of observations observed, and flight plan quality as measured by LOS WV for individual observations. Careful ordering is necessary to generate good quality flights.

Checking the constraints described above is complex, but not computationally intensive for a single flight. Unfortunately, solving a flight planning problem may require searching over very large numbers of flights. Consider the problem instance shown in Figure 2. Finding the best plan for this problem might lead to searching over *millions* of options.

Scheduling flights will also require choosing dead-legs to enable an observation. These might arise due to either SUAs that must be avoided, or to ensure the aircraft is at a place and time to satisfy the intersection of the observation duration constraints and elevation limits. Finally, scheduling flights will require selecting a good takeoff time for each flight. Both dead-legs and takeoff time selection require searching over continuous quantities, e.g. time and dead-leg headings. The number of schedules for such problems is technically infinite, and therefore it is practically impossible to search all schedules.

5. Challenges for Previous Approaches

Automated planning and scheduling has a long history. A complete survey is beyond the scope of this paper; however, we review some methods for automated planning and scheduling to see why they are unsuitable for automatically scheduling SOFIA flights.

A typical method of automatically solving planning and scheduling problems is to search the entire space of solutions, relying on a combination of heuristics to guide search towards good solutions very early, and efficient pruning techniques designed to eliminate infeasible or suboptimal solutions quickly. As we have argued above, the number of possible schedules (even without the problem of choosing start times or dead-legs) makes this approach unlikely to work.

Many scheduling problems assume that the earliest and latest times at which observations can be scheduled are provided up-front to the solver; these bounds can be adjusted to effectively prune schedules. These quantities may be calculated for astronomical targets observed when SOFIA is at a fixed location quite easily. If the aircraft is at position $\mathbf{p} = (\gamma, L)$, the earliest and latest times at which the observation is visible by SOFIA are given by the solutions to

$$\cos^{-1}((\sin(20) - (\sin\delta)(\sin\gamma))/((\cos\delta)(\cos\gamma))) + L + \alpha$$

The sin(20) term arises from the fact that SOFIA's lower elevation limit is 20 degrees. Depending on the time of year, location of the telescope and the target coordinates, some targets can be seen all night, and some cannot be seen at all. However, since the telescope is allowed to move, the earliest and latest times widen considerably. Further, there is now the problem of calculating the reachable positions of the aircraft, which depends on a variety of factors such as wind and weather. Finally, the upper elevation limit complicates the calculation of feasible observation times. Figure 4 shows the set of positions on the planet from which an observation can be seen. At any time, this set of positions is an annulus whose center is the position at which the target is directly overhead. This annulus rotates around the Earth; as described previously, the center of the annulus follows the latitude line at latitude δ, the target declination. Furthermore, for some targets, this geometry creates 2 visibility windows. Overall, the complexity of the constraints, looseness of the bounds, and computational costs of generating the bounds effectively preclude the use of pruning techniques.

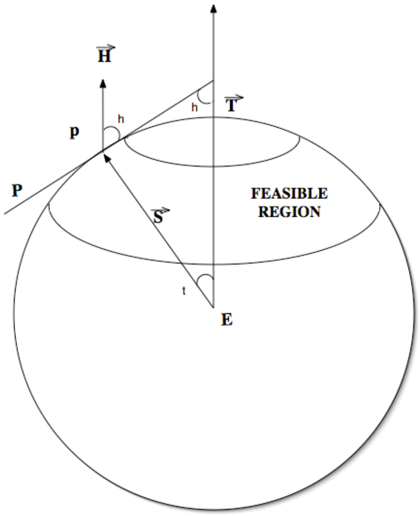

Figure 4. The set of positions on Earth from which an observation can be viewed at a fixed time.

A further complication arises because of the constraint that the aircraft takes off and lands at designated locations. Typically, the takeoff and landing airport are the same. The time and position at which an observation is performed influence the direction the aircraft flies. Observing a target while it is rising carries the aircraft South, since targets rise in the East and the telescope points out the left-hand side of the aircraft. Similarly, observing a target while it is setting carries the aircraft North. If the target declination is higher than the aircraft's latitude, the aircraft will be carried East; otherwise, the aircraft will be carried West. Careful scheduling is needed to ensure the aircraft lands at its destination while still ensuring that flights are of high quality.

6. Designing an Automated Flight Planner

In this section we discuss an approach to automated flight planning that addresses the challenges described in the previous sections. We begin with a server that provides an API for both automated and manual flight planning. The API makes use of a climatology model or weather prediction to provide wind, temperature and water vapor data. This API provides functionality to represent a flight planning problem, including the candidate set of observations and flight days; insert and remove observations from a flight, and calculate the implications; retrieve details from each leg such as target elevation, LOS WV, SUA incursions, fuel, outside air temperature, winds, and altitude.

We made the following assumptions for problem specifications: Target durations may not exceed 2 hours. A maximum takeoff fuel weight must be provided as part of the problem specification. The flight dates must be specified in advance, and the takeoff and landing airports must be designated for each flight date. The problem may specify calibration observations, and precedence constraints between observations and calibrations.

We made the following assumptions on solution methodology which limit the set of flight plans we can generate: The automated flight planner attempts to find a schedule maximizing the sum of TAC ranks of scheduled. An observation is either scheduled such that all relevant constraints are satisfied, or not scheduled; the automated flight planner does not attempt to divide observations into smaller pieces. Observations are guaranteed to be observed at a time at which the sun is below the horizon. Takeoff times are no earlier than half an hour before darkness at the takeoff airport, and landing times are no more than half an hour after sunrise at the landing airport. If a scheduled observation has a calibration, the calibration is scheduled on the same flight. Every flight obeys all relevant constraints on observation duration, telescope elevation limits, fuel consumption, and SUAs. A 20,000 lb fuel reserve is guaranteed for each flight. We chose a policy for constructing dead-legs that limits the set of flights we consider. If the aircraft must enable an observation by flying a dead-leg, we always choose the *shortest* dead-leg. This choice has three benefits: first, we do not have to search over an infinite number of possibilities; second, this leaves the maximum amount of time to observe targets later in the flight; third, it turns out that the search for the shortest dead-leg can be solved very fast; for details, see [2].

As previously discussed, complete search is unlikely to solve this problem. Previously we developed an approach based on constructing a flight from beginning to end. The approaches uses greedy search and sampling to decide which observation to add next. This approach, called ForwardPlanner, is described fully in [2,3]. Ultimately, this approach failed to perform well as more constraints were added to the problem; variants of this approach were fast, or produced high quality plans, but no approach did both.

We chose instead to use a form of algorithm known as *local search*. Briefly, these algorithms generate a candidate solution that might be either infeasible or have poor quality. The algorithms then generate modifications of the schedule in an effort to either reduce the number of violated constraints or improve the quality, and chooses one of these modifications as the new schedule. While unable to find the best solution, these methods often find high quality solutions very fast. We designed the AFP based on a local search algorithm called Squeaky Wheel Optimization (SWO) [7]. SWO takes as input a permutation of tasks to schedule, and schedules each task in the order specified by the permutation. Our version of SWO discards tasks if they can't be scheduled without violating constraints. The permutation and its resulting schedule are then analyzed to construct a new permutation that might schedule tasks that were previously rejected. The cycle repeats until all tasks are scheduled or for a fixed number of iterations. In Figure 5 we provide a sketch of SWO specialized to solve the flight planning problem. The first step is to generate a reasonable order for observations in the flight, then to decide on a takeoff time. Observations are then evaluated in order to see if they can be scheduled. If an observation is not trivially visible for the requested duration at the current position and time, the shortest dead-leg is found. If SUAs can be avoided, and sufficient fuel remains to return to the landing airport after the observation leg is completed, the observation is added to the schedule, otherwise it is rejected. All observations are processed in order; rejected observations are recorded.

Once this process is completed, SWO attempts to shoehorn the rejected observations into the existing schedule. This is done by considering each location and time at which an observation is performed in the schedule, and trying to schedule a rejected observation at that time and position. Generally, this may delay subsequent observations, causing them to be rejected.

The prospects for rejection of future observations and the relative TAC ranks of the observations involved are estimated and used to evaluate the net benefit of any reordering. This, in turn, is used to decide how to revise the ordering. Random sampling is used to ensure all rejects are given a fair chance in the next schedule. The best schedule of all generated is saved.

```
SWO(MaxFlights)
Generate observation order P
for MaxFlights
    Select the takeoff time θ
    for observation o in P
        if o can be performed at p,θ
            (possibly with a dead leg)
        and the aircraft can land afterwards
            Add o to the flight F
            Update p,θ
        else add o to reject list R
    Update the best flight B
    Modify(P,R,F)
end

Modify(P,R,F)
for observation o in R, each observation s in F
    $p_S$ is the position the aircraft before performing s
    $\theta_S$ the time the aircraft begins s
    if o can be performed at $p_S, \theta_S$
    and the aircraft can land afterwards
        v(o,s) is the net benefit of ordering o before s
        Add (o,s,(v(o,s)) to U
Randomly choose a reordering from U using
$P(o \text{ before } s) = v(o,s) / (\Sigma_{w(o,s) \text{ in } U}\ w(o,s))$
Modify and return P
```

Figure 5. The Automated Flight Planner algorithm, based on Squeaky Wheel Optimization (SWO).

Local search algorithms like SWO have the property that the more time they are allowed to run, the better quality results they provide. However, there is a "law of diminishing returns" that dictates how extra time leads to increased quality; performance tuning is necessary to get the best results. Furthermore, the entire process can also be repeated using a different initial ordering. We note that the algorithm sketched shows how to build a single flight; the generalization to series of flights is straightforward. For a set of representative problems, SWO can generate high quality flights quickly, providing a good combination of speed and quality. A comparison of SWO with ForwardPlanner is shown in Figure 6.

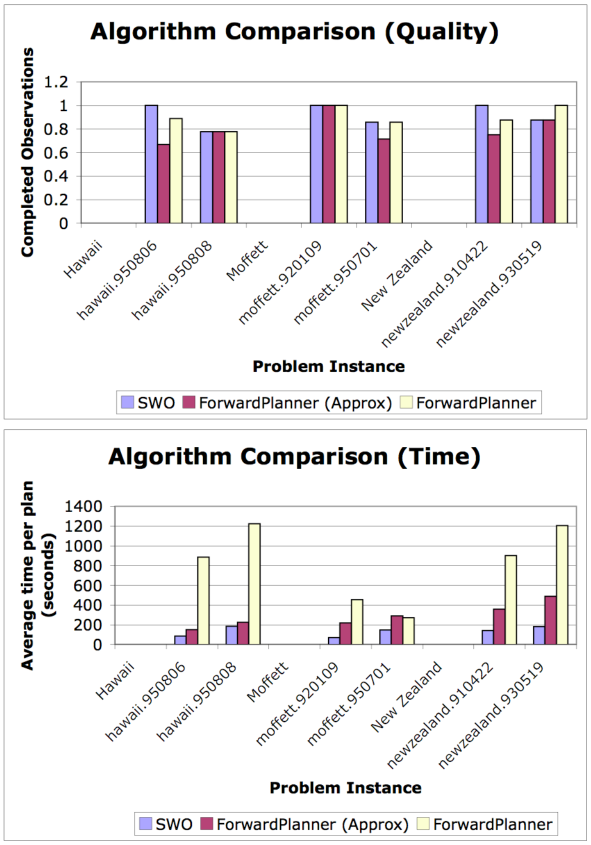

Figure 6. Quality (top) and time (bottom) performance of SWO on a small set of flight planning problems compared to previous approach.

Further details of the algorithm and comparison with previous approaches can be found in [4,5].

7. A New Challenge: Preferences

The automated flight planner described above was designed primarily to schedule as many high ranked observations as possible. However, there are other objectives that have arisen during the design of the planner; handling these objectives will require modifying the AFP to meet new challenges.

As stated above, a primary objective of SOFIA is to minimize the LOS WV achieved during observations. In [4] we describe a method for minimizing the LOS WV for an individual observation by forcing the observation to occur at the maximum telescope elevation. This can reduce LOS WV by a factor of 2-3, and dead-legs maximizing the telescope elevation can still be found quite efficiently. However, the optimization is "local" in the sense that each observation can be optimized alone, but may lead to poor choices for subsequent observations in the flight; in [4] we report that this approach leads to a tradeoff between achieving low LOS WV and scheduling more observations. Further work in this area is possible, by exploiting geographic variation in water vapor, as well as achieving aggressive altitude change by manipulating the initial fuel load.

Anecdotal evidence from astronomers who worked with KAO indicates that flight planners strove to ensure there were no dead-legs at high altitudes or near the end of flight plans. The automated flight planner cannot optimize flight plans this way. Some of the design decisions of the automated flight planner may even contribute to plans with dead-legs, especially at the end of the flight. Observing as early as possible may leave no suitable observation at the end of a flight, simply because all observations have been scheduled. As discussed in a previous section, observing targets while they are rising carries the aircraft south; the bias to observe targets as early as possible may lead the aircraft to fly south throughout most of the flight. With no impetus to observe targets while they are setting, the flight planner may be unable to delay observations and fly north. Finally, the flight planner is built in such a way that scheduled observations have a high likelihood of staying scheduled; scheduled observations are never re-ordered, and candidate reorderings of rejected observations that may "bump" scheduled observations are avoided. Searching for good flights will require relaxing these assumptions.

8. A New Challenge: Complex Tradeoffs

A highly automated flight planner will enable SOFIA flight planners to evaluate complex tradeoffs efficiently using less staff than was possible for KAO. In this section we describe some of these tradeoffs, and provide a demonstration of how the automated flight planner can help SOFIA meet these challenges.

If we return to the problem of scheduling flights with no dead-leg time, as described above, one way of doing this is to allow choice in the duration of observations, and express a *preference* function over the value of an observation of any particular duration. The automated flight planner can then search for flights with a quality measure that combines the TAC rank, scheduled observation duration, and dead-leg time. However, this would require astronomers to provide meaningful time-quality tradeoffs.

The approach outlined above assumes that a single quality measure can be assigned to a proposed schedule; however, sometimes this is not possible. There are complex tradeoffs between the takeoff fuel weight, flight duration, percentage of requested observations that can be performed, and average LOS WV for a flight. During flight, aircraft altitude is limited by aircraft weight. The more fuel carried, the longer the flight, but also the more limited the aircraft's initial operating altitude. Fuel is costly (JPA costing $3.00 a gallon in April of 2005, and even more today) so using it wisely is important. Furthermore, repositioning the aircraft to seek drier air or maximize telescope elevation will generally require longer dead-legs; this will reduce time for observing.

The constraints governing legal flights are complex enough that it is not possible to analyze the tradeoffs up front; one has to generate flight plans and compare them to see the tradeoffs manifested. Furthermore, the tradeoffs cannot be analyzed just once. The trade space will look different for different sets of observations, due to the complex nature of the aircraft's ground track. The trade space will also look different at different times of year; temperature and water vapor change throughout the year, affecting fuel consumption and LOS WV. Finally, the price of fuel is not constant, so operators' propensity to trade science for operations costs will also change. These tradeoffs are made even more complex because the quantities that are traded are usually incomparable. Lowering operations costs is valuable, but it is hard to put a dollar value on the science output of the entire observatory. Similarly, it is not possible (or at least quite difficult) to determine whether it is worthwhile to sacrifice one observation for lower LOS WV on another observation.

We demonstrate these tradeoffs in a study performed on the observations shown in Figure 2. All flights originate and terminate at Moffett Field, CA. Figure 7 shows the percentage of scheduled requests, fuel load and LOS WV of the schedules found by the AFP. We desire flights with low LOS WV and many scheduled observations, i.e. those in the upper left corner of Figure 7. We see that there are significant tradeoffs between takeoff fuel load, percentage of requested observations scheduled, and LOS WV. An extensive version of this analysis was originally published in [4].

In addition to desiring low LOS WV, astronomers may wish to impose constraints on LOS WV for sensitive observations. The automated planner described in [4] iteratively constrains the telescope elevation to reduce LOS WV; this is a promising way of managing these constraints.

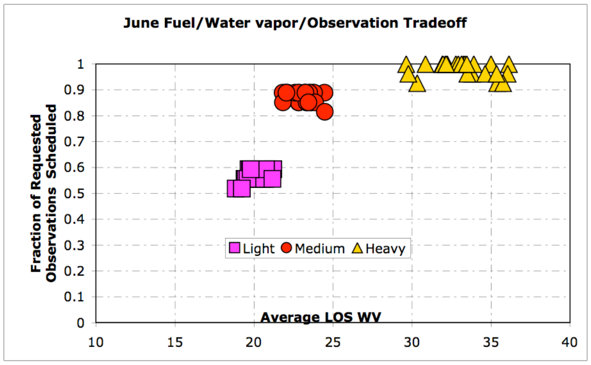

Figure 7. Tradeoff between takeoff fuel load, LOS WV and fraction of requested observations scheduled for June flight series.

9. SOFIA's Challenges: Planning Ahead

Automation can address many more of SOFIA's challenges. As mentioned, the automated flight planner assumes that requests specific to different instruments have been allocated to times of year before detailed planning of instrument blocks commences. Allocation of instrument blocks is presently done by hand, but the automated flight planner can potentially be used to allocate instrument blocks. A brute-force method of doing so is to try building schedules for each month of the year, initially trying to schedule every requested observation. Alternatives involve expanding on the basic SWO approach outlined for detailed flight scheduling. A more aggressive challenge, both computationally and operationally, is to use the AFP to mitigate operational uncertainty, e.g. due to weather. One strategy is to move automated flight planning onboard the aircraft in order to refine plans in flight. The operational challenge is to define the set of possible flight plan changes permitted by the observatory, astronomers, pilots, and the FAA. This challenge includes limitations of the onboard flight hardware and communications link to enable weather condition updates. The challenge is to formulate algorithms that will solve the appropriate rescheduling problems in order to preserve later observations.

10. Acknowledgements

We acknowledge the contributions of Elif Kurklu Michael A. K. Gross, William Vacca, Goran Sandell, Jacqueline Davidson, and Sean Casey over the years to this work.

11. References

[1] J. Bresina and A. Jonsson and P. Morris and K. Rajan, "Activity Planning for the Mars Exploration Rovers", *Proceedings of the 15th International Conference on Automated Planning and Scheduling*, Kluwer Academic Publishing, Monterey, CA, 2005, pp. 40-49.

[2] J. Frank and M. A. K. Gross and E. Kurklu, "SOFIA's Choice: An AI Approach for Scheduling Airborne Astronomy Observations", *Proceedings of the 16th Innovative Applications of Artificial Intelligence*, AAAI Press, San Jose, CA, 2004, pp. 828 - 835.

[3] J. Frank and E. Kurklu, "SOFIA's Choice: Scheduling Observations for an Airborne Observatory", *Proceedings of the 13th International Conference on Automated Planning and Scheduling*, AAAI Press, Trento, IT, 2003, pp. 226 - 235.

[4] J. Frank and E. Kurklu, "High and Dry: Trading Water Vapor, Fuel and Observing Time for Airborne Infrared Astronomy", *Proceedings of the 25th International Geoscience and Remote Sensing Symposium*, IEEE Press, Seoul, South Korea, 2005.

[5] J. Frank and E. Kurklu, "Mixed Discrete and Continuous Algorithms for Scheduling Airborne Astronomy Observations", *Proceedings of the 2nd International Conference on Constraint Programming, Artificial Intelligence, and Operations Research*, Springer Verlag, Prague, Czech Republic, 2005, pp. 183 - 200.

[6] M. Johnston and G.Miller, "SPIKE: Intelligent Scheduling of the Hubble Space Telescope", *In Intelligent Scheduling M. Zweben and M. Fox, eds*, Morgan Kauffmann Publishers, San Francisco, 1994, pp. 257 - 290.

[7] D. Joslin and D. Clements, "Squeaky Wheel Optimization", *Journal of Artificial Intelligence Research v 10*, AAAI Press, 1999, pp. 353-373.

[8] J. W. Potter and J. Gasch, "A Photo Album of Earth: Scheduling LANDSAT 7 Mission Daily Activies", *Proceedings of the International Symposium on Space Mission Operations and Ground Data Systems*, Tokyo, Japan, 1998.