JAMES WEBB
SPACE TELESCOPE

TECHNICAL REPORT

SPACE TELESCOPE SCIENCE INSTITUTE
Operated for NASA by AURA

Title: Monte Carlo Estimates of JWST Observatory Utilization	TM: JWST-STScI-000665[1] Date: 7 October 2005 Rev.: -
Authors: Knox S. Long Phone: 410-338-4862	Approvals: Peter Stockman

1.0 Abstract

The mission concept for planning and scheduling of JWST involves a sequence of related steps, which include proposal selection, long-range plan generation, and short-term scheduling. The ability to schedule effectively is determined by a variety of factors including the nature of the proposal pool, the constraints placed on scheduling by the architecture of the observatory, and the processes and algorithms that ultimately place observations in the short-term schedule. In this report, we use a set planning tools to perform a parametric study of parameters in a Monte Carlo generated proposal pool. These parameters – oversubscription, exposure lengths, fraction of orient constrained visits – are likely to affect our overall success in scheduling the Observatory, in light of requirements to keep momentum storage to an acceptable level on the Observatory. The study suggests that fixed orient observations are not a major constraining factor for limiting observatory utilization. The study also suggests that achieving a utilization of 97% would not require more sophisticated scheduling algorithms than employed here if other factors such as momentum management were not a problem. Momentum management is a problem though, and therefore we carry out a simple experiment to see whether limiting visit times of targets to those times when the magnitude of that torque is small significantly reduces the momentum management problem. We conclude that this is not an effective strategy for limiting the number of required momentum dumps on JWST. This may not be surprising since momentum is a vector quantity.

2.0 Introduction

The ease of scheduling observations with JWST will depend upon a variety of factors that fall into three categories:

a. The capabilities of the Observatory

b. The nature of the observing pool to be scheduled

c. The capabilities of the planning and scheduling system to plan observations on the observatory.

---

[1] This replaces STSCI-JWST-TM-000558, which should be considered obsolete.

In this memo, we explore some of the interactions between the nature of the observing pool and the capabilities of the Observatory in an attempt to shed light on desirable characteristics of the observing pool and scheduling strategies. In particular, we explore how effectively simple scheduling algorithms can be used to schedule JWST with observing pools that contain differing oversubscription rates and numbers of fixed-orient observations. We also conduct some preliminary studies relevant to the question of how momentum can be managed on the Observatory.

The study is based on Monte Carlo generated sets of potential visits using a set of software programs, which collectively goes under the name “planJ”. The software is described in Appendix A. Observations are only allowed to take place when they are in the field of regard of JWST, and at acceptable roll angles. Slew times are not calculated; instead an overhead time of 44 minutes is added to each observation to account for slew and setup times. In this memo, the term “visit time” means the total time that the visit would take, including any setups. The term “exposure time” refers to the time between the end of the acquisition and the end of the observation, i.e. this is time during which exposures could be made.

This study is described as a series of experiments.

3.0 Experiment 1 – Utilization as a function of visit length

Petro et al. (2004) showed that on-target times on HST, Chandra, and Spitzer follow a log normal distribution in time. They argue, based on the characteristics of JWST, that it was reasonable to expect that JWST exposure times would follow a log normal distribution with a mean exposure time of 10,000 s and a σ of 3x the mean exposure time. A second estimate of the JWST science program, the SO-DRM has an average visit time of 7,000s.

As a baseline for other studies, we have calculated the efficiencies for several sets of very simple observation pools with different mean exposure times, different angular distributions on the sky, different oversubscription rates, and different scheduling algorithms:

- Angular distribution: Isotropic, Galactic plane
- Mean exposure time: 20,000 s; 40,000 s; 80,000s
- Oversubscription: 1.0, 1.1, 1.2, 1.3
- Scheduling algorithm: Random/biased, Chronological, Fewest available targets, and Random/unbiased

The “isotropic” and “Galactic plane” distributions are intended to provide extremes of program pools in which a single region of the sky does not dominate the distribution. For the Galactic plane distribution, we have assumed that the observations are randomly distributed in longitude (L2) and are distributed in a Gaussian pattern in latitude (B2) with a σ of 3^{o}.

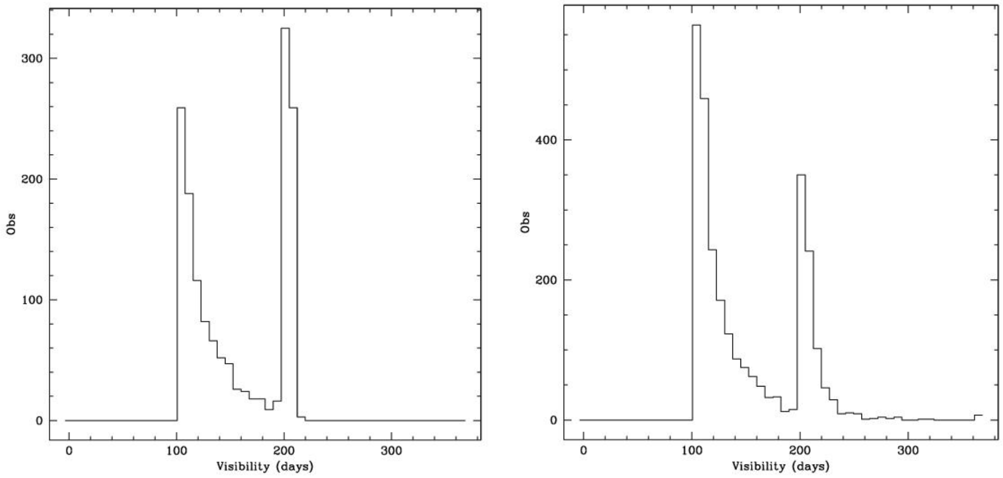

Figure 1—Visibility times of Galactic and isotropic target pools

The two figures above show the visibility distribution for a Galactic (left) and an isotropic target distribution (right). The average visibility is similar in both cases; 154 days for the Galactic plane distribution, and 145 days for isotropic distribution. On the other hand, the numbers of observations that are available for scheduling during the year is quite dissimilar, reflecting the rotation of the Galactic plane through the JWST field of regard.[2]

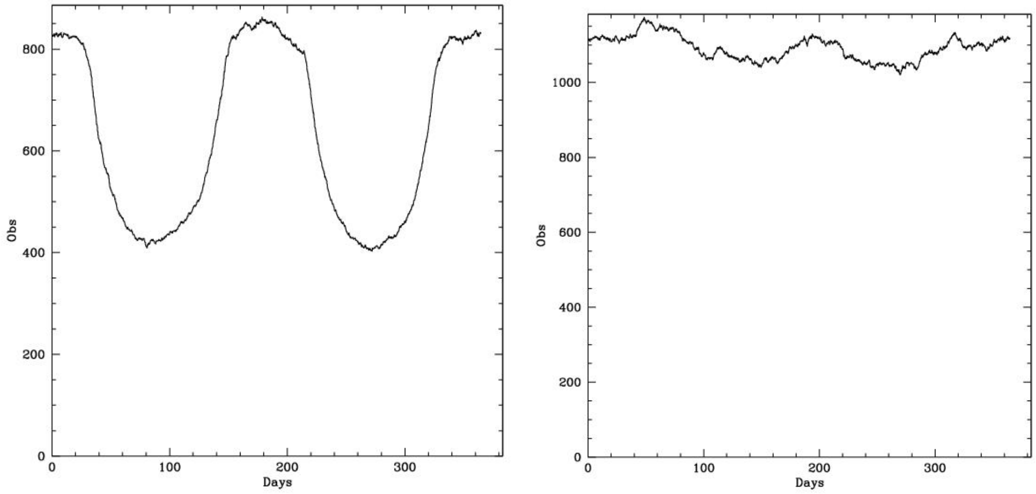

Figure 2 -- Oversubscription rates for Galactic and isotropic target pools.

The "planJ" software implements three basic scheduling approaches:

---

[2] The version of the planJ software being used for this study is the one that existed in June 2005. This version of the software does not include the capability to reduce scheduling windows below that determined by the field of regard and pitch-roll requirements. Therefore, it does not model the long range planning process, but rather attempts to create schedules from the full observing pool.

1. The first is a simple priority-based algorithm that schedules targets in a given order. Usually the order is something that is evaluated prior to scheduling the first target, such as the length of the observation. But the priority can be something that is evaluated periodically such as the ratio of observation length to current visibility for the target. The simplest example of a prioritization scheme is a totally random prioritization of the observations that is fixed at the beginning of the scheduling process.
2. The second basic scheduling approach is chronological. In this approach, observations are scheduled from the beginning of the scheduling interval, keeping the gap between the last observation scheduled and the next observation as small as possible. When more than one target can be scheduled at the same time, priorities are used to select which target to schedule.
3. The third approach is to find the place in the schedule with the fewest targets available for scheduling, and scheduling in that interval. Each time a target is scheduled, the information about the number of targets that can be scheduled is recalculated so that one continually attempts to schedule the time in the plan that has the fewest remaining targets. Since there will usually be several targets that can be scheduled, secondary criteria determine which of the available targets is actually scheduled in an interval.

In the “chronological” or “fewest available” approach, one selects the “time” first and then the target. In the “priority-based” schemes, the target is selected first and the “time” can be selected based on secondary criteria. There are two options that have been implemented with the priority-based scheme:

(a.) Schedule at the earliest opportunity (biased): In this case, the planJ software looks at all of the available times for a target and schedules it at the beginning of the first time it can be scheduled given the observations that have already been scheduled. This is intended to create a schedule that is tightly packed at the beginning of the schedule, since as many observations as possible are scheduled early on. Furthermore, since the visibility of a target being scheduled is most often determined by the targets that have been scheduled previously, this tends to produce a schedule without gaps, or at least without a lot of small gaps.

(b.) Schedule at a random opportunity (unbiased): In this case, the planJ software looks at all of the available times and randomly selects which time window to schedule into. For the first target, there will only be one or two visibility windows in the year, but the number of visibility windows grows considerably (to a number that is of order half the targets) as more and more targets are scheduled. In this approach, the planJ software randomly selects a visibility window in which to schedule the target and then randomly chooses to schedule the target at the beginning or end of that scheduling window. As is discussed later, this approach tends to sample the pitch angle distribution of the Observatory far more uniformly than the “earliest opportunity” approach.

The calculated utilization of the observatory for various oversubscription rates and scheduling algorithms is shown in Table 1. The specific approaches listed in the table are for (1) a random prioritization of targets to schedule placed as early as possible in the schedule (Random/biased), (2) scheduling the longest observation possible at the place in the schedule with the fewest targets (Fewest/longest), (3) scheduling from the beginning of the year selecting targets with the smallest ratio of remaining scheduling window to visit time (Chronological/ratio), and (4) a random prioritization of targets scheduling without biasing to the earliest scheduling opportunity (Random/unbiased).

Isotropic Distribution												
	Random/biased			Fewest/longest			Chronological/ratio			Random/unbiased		
Oversubscription	20000	40000	80000	20000	40000	80000	20000	40000	80000	20000	40000	80000
1.0	92%	93%	93%	94%	97%	95%	93%	92%	93%	99%	93%	96%
1.1	97%	95%	97%	100%	100%	100%	94%	97%	98%	99%	99%	99%
1.2	97%	98%	97%	100%	100%	100%	95%	97%	99%	99%	99%	99%
1.3	98%	98%	98%	100%	100%	100%	96%	97%	99%	99%	99%	99%
Galactic Plane Distribution												
	Random/Beginning			Fewest/longest			Chronological/ratio			Random/unbiased		
Oversubscription	20000	40000	80000	20000	40000	80000	20000	40000	80000	20000	40000	80000
1.0	90%	87%	92%	97%	98%	89%	91%	93%	92%	94%	95%	90%
1.1	97%	93%	96%	100%	100%	100%	95%	97%	95%	99%	98%	97%
1.2	98%	98%	98%	100%	100%	100%	95%	97%	98%	100%	99%	98%
1.3	98%	98%	98%	100%	100%	100%	95%	97%	99%	99%	99%	99%

Table 1 -- Utilization Efficiencies for target pools with no fixed slit observations

Generally speaking, all of the scheduling algorithms give relatively high utilization for an oversubscription rate of 1.1, of order 95% or greater. Multiple runs of the code were carried out for (most of) the cases explored. In each of the cases, a new “isotropic” or Galactic program file was created, and scheduled. Variations in utilization were small, of order 1% in nearly all cases.

A number of weak trends and patterns are evident:

1. The most efficient scheduling algorithm is the Fewest/longest - the one that schedules based on the number of targets that can be scheduled. This algorithm gives better than 99.5% utilization for a 1.1x oversubscription for all of the isotropic and Galactic samples. The “worst” algorithm appears to be the Chronological scheme, but even it gives 95-97% utilization at an oversubscription rate of 1.1
2. Average observation time has relatively little affect on the result of a particular oversubscription rate and algorithm. The parameter space we explored has exposure times longer than those of either the SO-DRM or the Petro et al. (2004) estimate based on Hubble, Chandra and Spitzer, but in view of the insensitivity to exposure time in the experiments we did conduct, it seems unlikely that our results will change if we were to go to shorter exposure times.
3. The distribution of targets on the sky has a negligible effect on the overall scheduling efficiency. At a subscription rate of 1, the isotropic distribution schedules a bit more efficiently than a Galactic distribution, presumably because the instantaneous subscription rate varies throughout the year for a Galactic distribution, as is shown in Figure 1.

In summary for this experiment, it is straightforward to schedule targets with average visibility periods approaching half a year.

4.0 Experiment 2 – Utilization as a function of fraction of fixed orient visits

Scheduling becomes increasingly difficult as constrained observations are added to the visit pool. One of the principal types of constraints is fixed orient visits. In the P&S operations concept, fixed orients arise either because there are scientific requirements to that demand an observations be taken at a fixed astronomical roll angle or because a fixed roll is established early in the planning process, as is likely to happen for nearly all multi-object spectroscopic observations with NIRSpec. Stockman (2005) suggests that the P&S should assume that orient requirements would be placed on 50% of the selected observation pool.

To test how fixed orients affect observatory utilization, we constructed schedules using observing pools with the following properties:

- Angular distribution: Isotropic
- Mean exposure time: 20,000s; 40,000 and 80,000s
- Percentage of observations with a fixed orient: 0, 20, 33, 50, 67, 100
- Oversubscription: 1.1

In the case of the planJ software, the orient angle for a fixed slit observation is assigned following the generation of the target position. The angle was assumed to be uniformly distributed among the allowed astronomical rolls for a target. Once the astronomical roll for the target had been established, the visibility windows were decreased to times when that astronomical roll was valid. The results of these experiments are shown in Table 2.

Isotropic												
	Random/biased			Fewest/longest			Chronological/ratio			Random/unbiased		
Fixed	20000	40000	80000	20000	40000	80000	20000	40000	80000	20000	40000	80000
0%	97%	96%	96%	100%	100%	99%	95%	97%	98%	99%	99%	98%
17%	96%	95%	95%	100%	100%	100%	95%	97%	98%	99%	99%	98%
33%	92%	93%	91%	100%	100%	100%	95%	97%	98%	99%	98%	96%
50%	92%	92%	88%	100%	100%	99%	95%	97%	98%	99%	96%	89%
67%	91%	89%	87%	99%	99%	99%	95%	97%	98%	98%	96%	90%
100%	91%	87%	84%	97%	98%	97%	95%	96%	96%	96%	92%	87%

Table 2 -- Utilization as a function of fixed orient fraction

The results of these tests were also encouraging. However, in this case, the scheduling approach that uses the number of targets that are available to schedule is clearly a bit better than the other algorithms. This is not surprising as it is the only algorithm that pays attention to the target pool as a whole.

4.1 Experiment 3 – Constraining visibilities by momentum build-up

The reaction wheels on JWST have limited capacity. The P&S operations concept for JWST has two major stages, the establishment of a long range plan and short-term scheduling. As demonstrated by Kinzel (2005), one way to manage momentum buildup is in short-term scheduling, where by adjusting roll angles and reordering observations one can significantly reduce the angular momentum build-up.

An alternative strategy is to reduce the scheduling windows for individual targets to those times when angular momentum build-up is low. If there were large variations in torque during the initial scheduling window, one would hope that limiting scheduling times to periods with relatively low torques would reduce momentum build-up.

To assess whether this might be an effective approach, we added a capability (in the "visibility" tool) to reduce the scheduling window based on the magnitude of the momentum build-up as a function of time. Specifically, if the visibility period in the absence of momentum considerations exceeds a user-specified maximum time, the tool reduces the scheduling windows to the maximum. This is intended to reflect the process in which visibility windows are reduced in long-range planning.

We then constructed multiple (9) schedules for isotropic and Galactic distributions of sources assuming a mean exposure times of 40 and 80 ksec and 1.1 times oversubscription using the same scheduling algorithms as previously. We compared the results when the plan windows were limited to 30 days and with those when the plan windows were not limited. As a measure of the effectiveness, we considered the 11-day rolling average rate of momentum accumulated per 11 days in the schedule (measured in N-m-s/day) and the utilization of the Observatory. (So that the comparisons of momentum build-up were fair, we renormalized to 100% utilization in all cases). The results are shown in Table 3. (Here we kept track both of the mean and the σ of the distributions for the 9 schedules, because there were some fairly significant differences in the <Mom> for the various runs.)

Isotropic40								
	30-days				*Unlimited*			
Approach	*<Util.>*	*σ (Util.)*	*<Mom>*	*σ (Mom)*	*<Util.>*	*σ (Util.)*	*<Mom>*	*σ (Mom)*
Random/biased	93%	1.2%	0.62	0.06	96%	2.2%	1.23	0.11
Fewest/longest	97%	2.1%	0.71	0.21	100%	0.8%	0.47	0.18
Chronological/ratio	97%	0.5%	0.66	0.08	97%	0.7%	1.26	0.11
Random/unbiased	95%	1.1%	0.21	0.01	99%	0.0%	0.14	0.05
Galactic40								
	30-days				*Unlimited*			
Approach	*<Util.>*	*σ (Util.)*	*<Mom>*	*σ (Mom)*	*<Util.>*	*σ (Util.)*	*<Mom>*	*σ (Mom)*
Random/biased	80%	1.9%	0.56	0.06	97%	1.1%	0.79	0.09
Fewest/longest	92%	1.3%	0.24	0.07	100%	0.0%	0.27	0.09
Chronological/ratio	83%	2.2%	0.53	0.10	97%	0.0%	0.65	0.15
Random/unbiased	83%	2.4%	0.23	0.03	98%	1.4%	0.20	0.07
Isotropic80								
	30-days				*Unlimited*			
Approach	*<Util.>*	*σ (Util.)*	*<Mom>*	*σ (Mom)*	*<Util.>*	*σ (Util.)*	*<Mom>*	*σ (Mom)*
Random/biased	91%	1.9%	0.56	0.10	95%	1.3%	1.28	0.23
Fewest/longest	97%	1.7%	0.78	0.26	100%	0.4%	0.31	0.19
Chronological/ratio	95%	2.2%	0.51	0.07	98%	0.3%	1.13	0.13
Random/unbiased	90%	1.9%	0.24	0.05	98%	0.4%	0.22	0.13
Galactic80								
	30-days				*Unlimited*			
Approach	*<Util.>*	*σ (Util.)*	*<Mom>*	*σ (Mom)*	*<Util.>*	*σ (Util.)*	*<Mom>*	*σ (Mom)*
Random/biased	79%	2.6%	0.42	0.07	96%	0.9%	0.67	0.19

Fewest/longest	91%	2.2%	0.31	0.17	100%	0.3%	0.40	0.34
Chronological/ratio	83%	3.9%	0.51	0.29	97%	2.8%	0.73	0.25
Random/unbiased	81%	1.6%	0.25	0.06	97%	1.3%	0.26	0.11

Table 3 -- Effect of limiting visibilities using torque magnitude. The momentum columns show the momentum accumulated per day averaged over rolling elcven day intervals.

The first thing that is obvious from Table 3 is that when the full scheduling windows are utilized there are considerable differences in the momentum build-up rates depending on the scheduling approach. Using a scheduling approach that schedules targets in random order but at the beginning of their scheduling windows (Random/biased) is far worse from a momentum perspective than the approach that schedules in random order, but "slides" half the targets forward to the beginning and half the targets to the end of their scheduling windows (Random/unbiased). Indeed, randomly choosing targets and randomly placing them at the beginning or end of their scheduling windows seems to be the best of the scheduling algorithms for both full visibility and reduced visibility windows. The location where targets are scheduled in the instantaneous field of regard is shown in Figure 3 for the Random/biased and Random/unbiased scheduling algorithms. This dramatically illustrates the asymmetry of the time-biased scheduler, which translates into poor balancing of angular momentum. Scheduling in chronological order is also relatively poor from a momentum perspective, presumably for similar reasons.

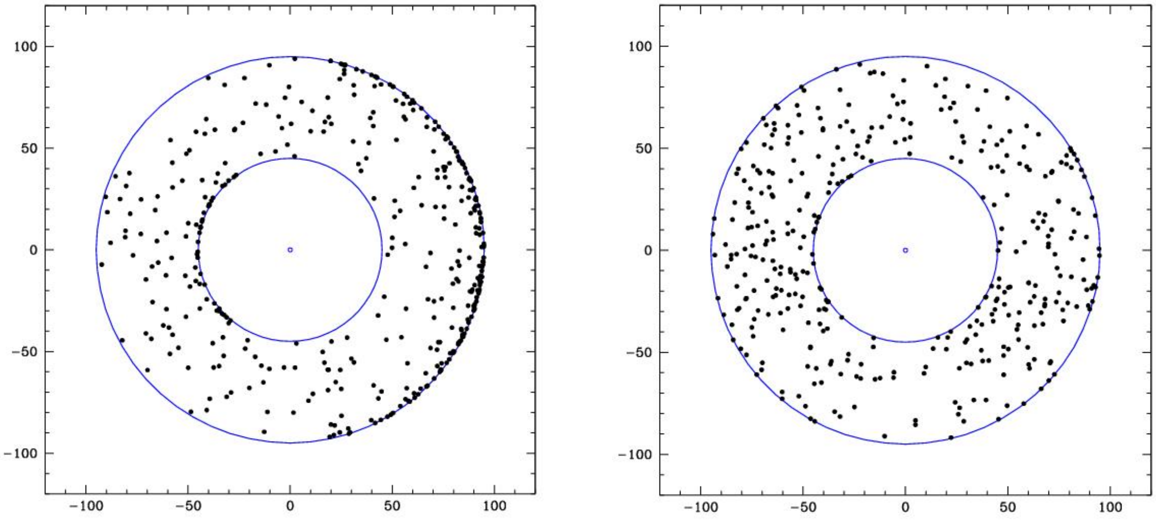

Figure 3 -- Where targets are schedule in the instantaneous field of regard for the two scheduling algorithms that choose and schedule targets in a random order. A given target enters the field of regard from the right and moves horizontally across the plot at a fixed ecliptic declination. The left hand panel shows the result when targets are "slid" forward to the earliest possible time (Random/biased); the right panel shows the result when there is and equal probability of sliding targets forward and back (Random/unbiased).

Reducing the scheduling windows to the lowest torque 30 days reduces the utilization that is achieved with the scheduler. The magnitude of the reduction in utilization varied, with the largest reduction for the two "random" algorithms, where for the 80-ksec Galactic source distribution the utilization for the Random/unbiased algorithm dropped from 97% to 81%. This is not entirely surprising since there are fewer opportunities to schedule with smaller scheduling windows. The amount of reduction is smallest for the "Fewest available" algorithm, a result that is not entirely surprising since this is the only algorithm that applies planning concepts by reanalyzing what targets to schedule given the targets available at any time. Reducing the scheduling windows reduces the momentum buildup for the algorithms that had the worst momentum properties, i.e. for the "Random/biased" and the "Chronological" algorithms. It did not help for the other "Random/unbiased" and "Fewest available" algorithms that had reasonably good momentum performance with the full visibility windows. Indeed all of the scheduling algorithms seem to produce rather similar performance with respect to momentum management once the scheduling windows are reduced. As illustrated in Figure 4, the distribution of target positions in the instantaneous field of regard looks much more similar for the random/biased and random/unbiased scheduling algorithms with reduced scheduling windows.

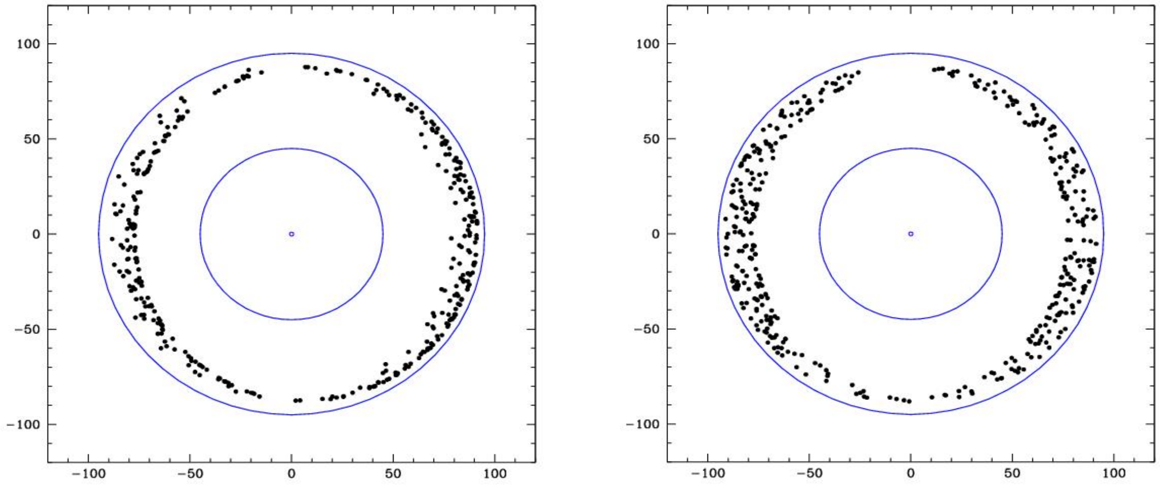

Figure 4 -- Position in instantaneous field of regard when maximum scheduling window is reduced to 30 days for a schedule created from an isotropic distribution of targets with mean exposure time of 80 ksec from the "random/biased" (left), and "random/unbiased" (right) algorithms.

Based on this experiment, we would conclude that limiting scheduling windows using the magnitude of momentum build-up does not reduce momentum buildup significantly, unless the algorithm chosen had very poor performance initially. To make progress in this regard, one almost surely has to take the vector nature of angular momentum buildup directly into account.

5.0 Summary

In this Technical Report, we have described a series of simple experiments based on Monte Carlo generated sets of observations and several very simple algorithms for scheduling the observations. Any conclusions drawn from the study have to be taken as preliminary. However, if these were the actual scheduling algorithms used then we would have to conclude the following:

1. Oversubscriptions of 1.1 are sufficient to provide a utilization rate of 97% for JWST using the extremely simple set of algorithms defined here.

2. The fraction of fixed slit observations is not in and of itself a strong determinant of scheduling efficiency, until the number of such observations dominates the program. (This is not to say that fixed slit observations are not more difficult, but it does suggest that the difficulty arises from other factors than the mere ability to place them in the schedule.) Of the scheduling algorithms that currently exist in the planJ suite the algorithm that schedules the portion of the year with the fewest available targets is the best for maintaining high utilization with low oversubscription and a high fraction of fixed slit observations.
3. Using the magnitude of the angular momentum build-up is not a sufficient momentum management strategy on its own.

6.0 References

Kinzel, W. 2005, STSCI-JWST-TM 000540, Managing Angular Momentum Accumulation by Visit Sequencing and Visit Roll Selection

Petro L, & the SO-DRM Working Group 2004, STSCI-JWST-TM-2004-0011, Science utilization patterns of JWST

Stockman P. 2005, JWST-STScI-000490, Implementation of NASA Science Policies related to JWST Planning and Operations

7.0 Appendix A – A description of the “planJ” Software package

The planJ S/W package is a series of tools, written in c, that emulate certain aspects of proposals selection, planning and scheduling on JWST. They have been tested on Linux, Solaris, OSX, and Windows (Cygwin & Microsoft c). The tools are continuing to evolve and anyone interested in the S/W should contact the Knox Long (long@stsci.edu). The tools are driven by ASCII parameter files, and the inputs and outputs to the various routines are all in ASCII (keyword value) formats. At June 2005, the S/W consisted of the following tools:

7.1 Pretend

The routine “pretend” expands a set of “programs” to a set of observations or “visits”. Each program is defined by an angular distribution of the targets on the sky, the mean visit time (which can be randomized to a log normal distribution in the output observations), the instrument, and whether the observations require a fixed orient.

Unlike JMS, which contains a model of each instrument, and expands exposures to construct the length of a visit, the user specifies the visit length in the inputs to the planJ software. The user also defines a fixed overhead time that is intended to account for the mean slew and setup time; the overhead time is fixed for all observations.

Finally, the user specifies the duration of the observing pool to be generated and optionally the amount of oversubscription desired for the observing interval. The program expands or contracts the observing pool in order to provide the desired oversubscription rate.

“Pretend” is a Monte Carlo code, and normally generates a different set of observations each time it is run.

7.2 Visibility

The routine “visibility” determines the visibility intervals for the observations, given the pitch and role constraints imposed by the JWST architecture. It also assigns fixed orients to those observations that require them. One can optionally limit the scheduling windows to times in which the momentum buildup is less than a certain amount.

7.3 Schedule

The routine “schedule” attempts to schedule the observations given their visibilities. There are 3 basic scheduling algorithms, as was described in Section 3.0. The three algorithms are based on a priority scheme, based on scheduling chronologically, or based on the number of targets that can be scheduled at any given time. The algorithms can be combined in various ways, so that for example one can choose the target with the longest visit time that can be scheduled at the time in the year where the fewest targets are in the plan.

All of the scheduling approaches incorporate features where the scheduler tries to slide observations forward or backward in time to fill in gaps. Once an observation is scheduled it is not moved in the schedule after that. (It does not slide one observation to create a gap for another observation, nor do two or more scheduled observations slide to make room for another observation. The sliding feature creates more efficient schedules but also ones in which most observations occur close the beginning of their visibility windows and at the most extreme pitch angles.)

7.4 Verify

The routine “verify” checks that visibility constraints have not been violated, calculates the momentum buildup for the observations from a set of torque tables, and provides certain other summary information about the schedule. Observations with fixed roll angles are scheduled at those roll angles; observations which are not roll-constrained are assumed to be scheduled at 0 roll, which minimizes the momentum buildup.