IRS AOT Revisions Spitzer Gyro Performance

C. Grillmair & H. Teplitz

IRS Instrument Support Team, Spitzer Science Center

07/23/04

I. Background

Two separate runs of the IRS Large Offset Test as well as independent measurements by the IRS GTO team have demonstrated that, given the current performance of the gyros, it is not possible to execute a significant number of existing IRS AORs without severely compromising the science (ie. partially or completely missing the targets). Despite considerable ongoing effort by the OET and the PCS working group, it has thus far not been possible to identify the root cause of the problem. While in-flight experiments are planned to better characterize gyro performance and to test a hypothesis, the IRS instrument support team has been asked to examine procedural changes or work arounds which would preserve IRS science in the event that hardware performance cannot or will not be improved.

The scope of any work around depends both qualitatively and quantitatively on the nature of the problem. An improved bias estimation filter by OET may help to reduce the frequency and magnitude of poor bias estimates, and thus the incidence of high drift rates. However, it is not yet known whether such an improvement is possible, or to what extent it will solve the problem. If the problem stems at least partly from a poor A2X scale factor (as has been suggested from time to time), then the least expensive, most prudent solution is to improve the scale factor on board the spacecraft, rather than modify command sequences on the ground to cope with the problem. Neither of these hypothesis is completely consistent with the available data (the maximum bias error quote by J. Tietz is only 1/3 to ½ of the amount necessary to account for the pointing drifts seen in the IRS large offset test, but a simple error in the A2X scale factor is not consistent with all the data either).

II. Case 1: Noisy Gyro Bias Estimates

If our pointing problems are primarily the product of inaccurate bias estimates *and* the solution is not to update the bias estimate for extended periods of time, then the work around for IRS is not costly, and indeed more efficient from an observer's point of view. All IRS high accuracy modes invoke bias estimation (a 120 second inertial hold followed by a PCS_GYRO_UPDATE) every 15 minutes.   Medium and low accuracy AOTs do not invoke internal bias estimates – these are inserted into the schedule between AORs.

If the pointing problem can be identified with unreliable bias updates, and a less frequent but more well measured bias estimate is the cure, then the changes to IRS AORs would be limited to simply removing all inertial holds and PCS_GYRO_UPDATEs. A separate (and welcome) action would fall on OPST to delete once-per-three hour invocations of gyro updates from standard scheduling protocol.

Recoding of AIRE would take XX days. There is a remote possibility that the fix could be incorporated into the S11 delivery. Integration and test would then be accomplished with little additional cost. However, if we have not identified the source of the pointing problem in the next few weeks, it is unlikely that the fix could be included in the S11 build schedule. The next scheduled delivery would be S12, with full-up implementation by Cycle 3. It is not reasonable to keep a large number of high accuracy or cluster AORs on hold for more than a year so we would need to schedule a point build to accommodate the fix.

III. Case 2: Pointing Issues Remain Unresolved or Uncorrected

This situation has more serious consequences for observing efficiency and development effort.

For a high accuracy staring AOR, a single 60 second exposure time (cycles=1) and INC_POINT_B, we find that after the first dither position on the Short-Lo-1 slit:

Sigma = SQRT(0.28”^2 + 0.2”^2)+90s *.004”/s = .704” single axis.

Here 0.28” is the single axis uncertainty immediately after a peak-up, 0.2” is the “bump” contributed by star tracker bias in the course of every move, 90 seconds is the amount of time spent slewing and then holding on gyros in INC_POINT_B, and 0.004”/s is the magnitude of the uncorrected gyro drift observed in IRS Campaign 9. (Note that this is about half the magnitude of the drift we saw in IRS Campaign 4). Slew times for most AORs are anywhere from 5 to 60 seconds, and we take 10 seconds as a typical time to slew from one slit to another. INC_POINT handoff times are always 80 seconds. The gyro drift is not added in quadrature since its direction does not appear to be random over the course of a single AOR (indeed, the direction is remarkably constant over many successive bias estimates). We see from the above that the pointing uncertainty has spanned the breadth of the HardPoint1 window in a single move from the peak-up array to the first requested slit.

Propagating the uncertainties over the standard sequences of dithers and slits, we find:

FOV	Sigma (arcseconds single axis)
SL1-b	1.1
SL2-a	1.5

SL2-b	1.8
SH-a	2.2
SH-b	2.6

We are well beyond even a QuickPoint condition in the course of a single, standard, multislit, staring AOR.

One might argue that adding the gyro drift in this way isn't entirely appropriate. Let us assume that we can't predict the direction of the drift, but that the magnitude will be around 15"/hr = .004"/sec (e.g. IRS Campaign 9). If the drift is along the slit, it is likely to be inconsequential for science. The situation is more serious if the drift is across the slit. If we adopt 1/sqrt (2)*0.004"/sec = .0028"/sec, then the above table becomes:

FOV	Sigma (arcseconds single axis)
SL1-a	0.6
SL1-b	0.9
SL2-a	1.2
SL2-b	1.4
SH-a	1.7
SH-b	2.0

In both cases we are well beyond the upper bound of the high accuracy mode. If we assume the gyro drift case of 4 mas/sec, then we would need to do a peak-up before every dither position to meet the .71 arcsecond Hardpoint-1 ceiling. Given that peak-ups can take up to 3 minutes per invocation, this option does not warrant further discussion.

From a command work around perspective, there are two more viable solutions:

(1) Frequent Attitude Resets

We can insert PCS_ATT_RESETs before or after every telescope movement (i.e. change of target and/or change of dither position or slit). The costs associated with this solution include:

a. A 20 second hold before each attitude reset.
b. An approximate doubling of the number of telescope moves for a given number of targets.
c. Pointing drift over the course of each exposure.
d. Development, integration, and test as in (II) above.

The time penalty in (a) could in principle be ameliorated by placing the attitude resets *after* each exposure so that the 20 required inertial seconds become part of the exposure time. However, since pointing errors in INC_POINT_B will depend at least partly on the length of a slew (whether due to poor bias estimates or incorrect scale factors), attitude resets might be better carried out after the slew and before the exposure.

If we put an attitude reset *before* every subslit position, then at the beginning of each subslit exposure:

Sigma = SQRT (.28^2+0.2^2) = .34 arcseconds single axis

To within the uncertainties, this method appears more than adequate to maintain sigma < 0.71 arcseconds.. For an all-slit, single cycle exposure sequence with 60 second exposures, the cost of these pre-exposure attitude resets would be 6*20=120 seconds, or about a 20% increase in execution time. If we include the initial slew to the target and the initial peakup, the cost is reduced to less than a 10% increase in execution time per AOR. The percentage cost would be reduced for multiple exposure cycles at each pointing.

If we put an attitude reset *after* every exposure (in an effort to save up to 20 seconds of settle time per move), then the pointing uncertainty at the beginning of every exposure will be:

Sigma = SQRT (.28^2 +0.2^2) + 90.*.0028 = 0.6 arcseconds single axis

for 60 second exposures and our most optimistic characterization of the gyro drift above. For most applications this represents effectively an upper limit, since for shorter exposures the time we spend on the drifting gyros is reduced and the second term becomes negligible, and for longer exposures we hand off to the star tracker, never exceeding about 90 seconds on the gyros (although very long slews on the order of one degree can increase the time on gyros to 140 seconds, yielding a sigma of 0.74 arcseconds). We can therefore just maintain Hardpoint-1 in a high accuracy staring AOR without invoking additional peak-ups. The efficiency cost would be greatest for the shortest exposures, where only ~7 seconds of the required 20 second settle time could be absorbed into the exposure itself, thus requiring an additional 13 second delay before moving on to the next target or subslit. For all other exposure times the extra settling time would be shorter than the DCE time, and observing efficiency would be unaffected.

While we still do not understand the underlying reasons for the occasionally poor gyro performance, we have now identified several data sets showing wildly out-of-family gyro performance. If we assume these data sets to (a) represent “worst case” abnormal behavior, and (b) be entirely a product of gyro drift, it may be possible to reduce the frequency of attitude resets, and thus the penalty to observing efficiency. If we adopt 0.007 arcseconds/second as our highest inferred gyro drift (IRS Campaign 4), then to maintain HardPoint-1 we could run purely on gyros for an accumulated total of no more than 50 seconds before requiring another attitude reset.

We conclude that, if we choose to maintain the best possible pointing accuracy using frequent attitude resets, the option which best serves this purpose is to invoke the attitude resets prior each exposure. The disadvantage of this approach is that it requires an additional 20 second settle time which must precede each attitude reset.

- Note (1): If the Observer is well settled at the time of the attitude reset, we could choose to reset the attitude using the Observer rather than the Star Tracker, thereby reducing the 0.2" STA error term in the sigma computation to ~0.02 arcseconds. The settling time required is the subject of an IRAC pointing test scheduled in early September.
- Note (2): Since in INC_POINT_B we are still relying on the gyros for the first 80 seconds at each pointing, the pointing during an exposure may drift by up to 0.2" (0.7" in the worst case).

(2) Observer_Always

The second work around available to us in the event that gyro performance cannot be improved is to routinely operate using an Observer rather than an Incremental Pointing mode. The current procedure for almost all maneuvering within AORs is to use INC_POINT_B, which amounts to running purely on gyros (IRU_ONLY) for the slews between targets or apertures, remaining under gyro control for an additional 80 seconds, and then handing off to Observer_B. If gyro drift cannot be adequately measured or corrected on board the spacecraft, then we could avoid running solely under gyro control by staying on one of the Observers.

The advantage of being purely under gyro control is that it provides for the shortest possible settle times. If we choose to slew and settle using an Observer, then we must wait longer before starting an exposure. Using SPOTIER 9.0.2 and the current version of the slew model, we can compute the approximate differences in slew+settle time per target using the Pointing Ready Indicator (PRI) (though OET has indicated that the PRI is rather conservative, and that spacecraft settles considerably faster PRI predictions). We assume that the current IRS practice of overlapping slew/settle time and detector conditioning time ("Efficient Step-and-Stare") would be carried over to an Observer_Always mode. Figure 1 shows the time increase per target as a function of slew distance that would result in going from INC_POINT_B to PCS_OBS_B with a PRI sigma of 0.4 arcseconds (Hardpoint-1). The figure demonstrates that for exposures of 30 seconds or less, PCS_OBS_B would cost an additional 10.5 seconds of settle time per target. On the scale of the focal plane (10 arcminutes or more), for exposures of 120 seconds or longer there would be essentially no difference between INC_POINT_B and PCS_OBS_B, since the time from the beginning of the slew to the beginning of the exposure is determined entirely by the time required for detector conditioning.

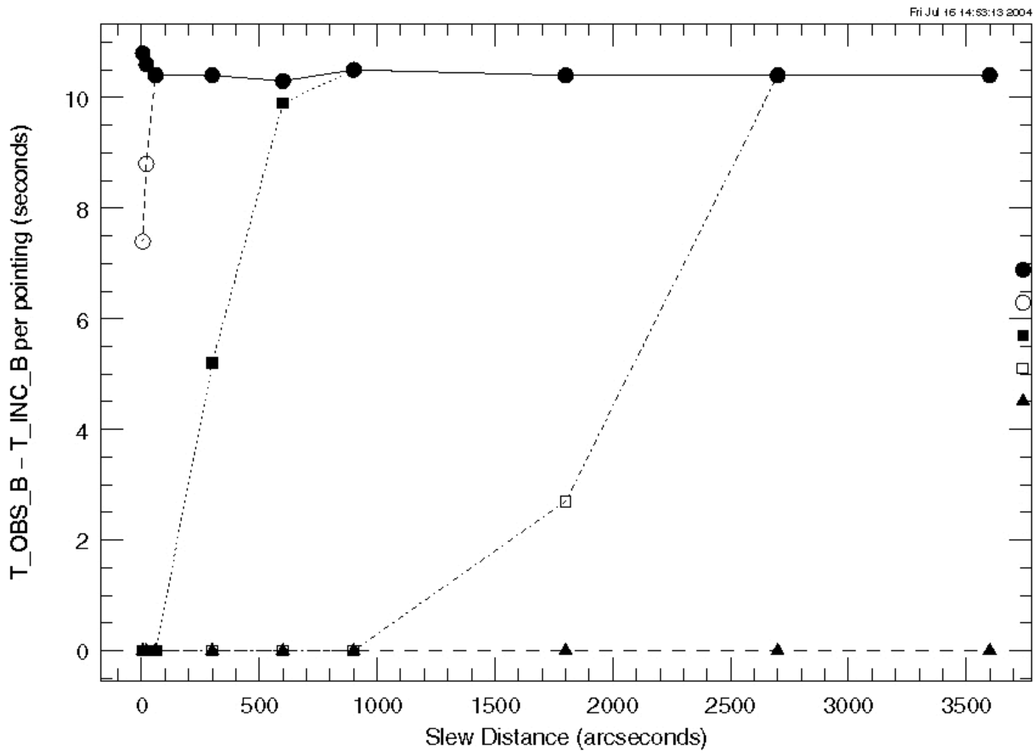

Figure 1. IRS AOT time increase per target that would result in going from INC_POINT_B to PCS_OBS_B as a function of slew distance. The symbols correspond to exposure times as follows: filled circles – 6 seconds, open circles – 14 and 30 seconds, filled squares – 60 seconds, open squares – 120 seconds, filled triangles – 240 and 480 seconds.

If we increase the allowable PRI sigma to 1 arcsecond (corresponding to Hardpoint-2), the settle time required is reduced to 10.0 seconds beyond what we are currently getting with INC_POINT_B. Interestingly, if we switch PCS_ATT_DET to Observer_A, the total slew+settle times are identical to those measured using Observer_B, even though Observer_A is supposed to settle more quickly.

Assuming that there are no other consequences to switching to an Observer-Always mode (the IRAC team will be testing such a strategy in September), how badly would such a change affect IRS efficiency? Due to all the possible options and modes available within the IRS AOT, it is difficult to estimate precisely the consequences of such a change without a detailed analysis of all AORs currently in the Spitzer database. However, we can get a rough idea by examining a subset of IRS AORs. 93% of IRS Staring AORs using Short-Lo1 in the SODB use fixed, single targets, accounting for 89% of the time spent using this AOT/aperture combination. Using the extra settle times indicated in Figure 1, multiplying by 2 to account for the two dither positions for each target, and multiplying by 6 to estimate the contributions of all other slits, we arrive at a total of

302,000 seconds. This is about 3.5 days out of about one year's worth of IRS observations, or an efficiency reduction of about 2.4%.

IRS mapping AORs are another matter. While large area mapping generally does not require great accuracy, a gyro drift of 9 mas/sec could result in significantly larger or smaller grid point spacings. In the special case where an observer needs to ensure complete areal coverage of a target, he will need to adopt a worst case view (that the drift will be high and perpendicular to the long axis of the slit) and shrink his map spacings accordingly. For 60 second integrations in the short-lo slits, a 9 mas/sec gyro drift could lead to shortfalls (or overshoots) of about 1 arcsecond per pointing. From a nominally optimal overlap of about 1 arcsecond (to compensate for STA bias at the 3-sigma level), one would need to increase the overlap by an additional arcsecond, reducing mapping efficiency with Short-Lo by ~60%. In this instance the observer would benefit from an Observer-Always mode, since the time to slew and settle in PCS_OBS_B is shorter than the detector conditioning time. However, someone doing large area, short exposure maps will incur a 10 second penalty per grid point, potentially adding many minutes to each map, or cutting down the total area coverage by up to 60%.

Overall, this option appears preferable to invoking frequent attitude resets. The increased settling time is roughly halved, there would be no pointing drift while exposing, reliance on gyro attitude estimates would be minimized, and as long as sufficient settling time is provided, the pointing would never degrade below the star tracker calibration.

Is there any science that can't be done at all using attitude resets or Observer_Always? I can't think of anything.

IV. Conclusions

If the current pointing difficulties are resolved and implemented entirely on the spacecraft, no changes to current IRS command sequences would be required.

If the pointing problems are determined to be a product of noisy gyro bias estimation, it may be necessary to remove initial and intermediate gyro bias estimates from high accuracy IRS AOTs, and from the OPST scheduling protocol. While some SSC development and test effort would be required, such changes would have positive consequences for observing efficiency, ranging from 0 to 10%.

If the pointing problems are not resolved, from the standpoint of both pointing accuracy and observing efficiency, the most attractive option would be to switch to an Observer for all incremental telescope pointing. This would maintain pointing at star tracker levels, and (if the PRI is to be believed) introduce no more than 10.5 seconds of additional settle time per pointing. This would require a considerable development and test effort prior to implementation, and would cut IRS observing efficiency by ~2.4%.