

# **Wide-field Infrared Survey Explorer (WISE)**

## **Frame Co-addition Peer Review Report**

**24-January-2008**

**Prepared by: Roc Cutri & Frank Masci**



**Infrared Processing and Analysis Center  
California Institute of Technology**

WSDC D-A001

## **1 INTRODUCTION**

A peer review of the design of the Frame Coaddition module of the WISE Science Data System was conducted by the WISE Science Data Center on November 15, 2007. The Frame Coaddition module, (AWAIC – A WISE Astronomical Image Coadder) is responsible for combining and mosaicing multiple WISE image frames onto a common pixel grid by reprojecting, resampling, undistorting and optimally combining pixel values. The module incorporates throughput scaling, background matching, outlier rejection and uncertainty validation.

The functionality, design and algorithms for the AWAIC module are described in the Image Coadder/Mosaicer (AWAIC) Subsystem Design Document (WSDC D-D005).

### **1.1 Review Panel Members**

Sean Carey (IPAC/SSC)  
Nick Gautier (JPL/WISE Science Team)  
John Good (IPAC/IRSA)  
Russ Laher (IPAC/SSC)  
Dave Shupe (IPAC/NHSC)

### **1.2 Instructions for Review Panel**

The peer review panel was asked to comment on the following specific questions:

- Does the design of Frame Co-addition program address the requirements on the system?
- Are the Frame Co-addition algorithms suitable and appropriate to carry out the system functions?
- Is the design robust to circumstances that will be encountered with the WISE data?

In addition, comments on other aspects of the design were welcomed.

Written reports were received from Carey, Good, Laher and Shupe. These reports are provided below. Miscellaneous comments that were recorded by R. Cutri during the review are also summarized below.

### **1.3 Applicable Documents**

WSDC Functional Requirements Document (WSDC D-R001)  
WSDS Functional Design Document (WSDC D-D001)  
Image Coadder/Mosaicker (AWAIC) Subsystem Design Document. (WSDC D-D005)  
WISE Calibration Plan Document (JPL D-33753)

## 2 PANEL REPORTS

### 2.1 *Sean Carey*

For outlier rejection, using the temporal rejection as Spitzer does should work more than adequately as you have plenty of samples (minimum of 8). Using mad to calculate the sigma is definitely the way to go if the outliers occur on one side of the distribution more than the other (radhits will be positive). One cautionary note is not to regrid the data to too fine a mesh when stacking the input pixels as the interpolation can de-weight the significance of outliers. This may influence your implementation as Frank has proposed a one pass outlier-mosaic scheme which is very I/O efficient. However, if your output grid is too fine, outliers may leak in more than you would like. Gridding the data on at lower resolution for the outlier detection might be a better option(you'll need some data to try it on).

In any event, you may want to consider "growing" the outlier masks to catch the edges of the radhits. For IRAC, the Si:As arrays tend to have fatter radhits (as the detectors are physically thicker). For these arrays, I used a low fraction of radhit overlap to output pixel overlap as a threshold. It's conservative, but that's okay as you have plenty of samples. While I am thinking about it, you will not want to apply the image construction kernel weighting for the outlier detection gridding and use straight area overlap instead.

Most asteroids will probably only be partially rejected as you won't use a low enough sigma threshold to reject the wings of the PSF for the asteroids. That's okay, but you might want to have a flag in the QA or post-processing stage to indicate that this is happening. The cores of the asteroids should be well rejected though as they will be significant positive outliers in 50% or less of the coverage. That's what we noted in MIPS GAL in any event.

Background matching for large maps is a challenge. One thing you will have to decide upon is whether you are going to remove a model of the Zodiacal light or not. If you are going to provide a product that could be used for extended structure studies, I would recommend considering zody removal (especially if you get the second epoch data!).

I'm attaching a draft of the MIPS GAL 24 um processing paper by Don Mizuno. Section 5 has an extensive discussion of our application of the overlap correction. You could implement a similar solution using the sparse matrix direct solution that Don did or a relaxation method as I think Montage does. The important thing in the MIPS GAL background solution is the damping terms applied. Please do not distribute this paper as we have not yet submitted it.

Be careful, if you perform a bootstrap overlap correction method where you start with a one pair of images and propagate the solution (my understanding of the 2MASS method). This type of solution can be very unstable over large solid angles.

I haven't quite convinced myself that PRF weighting is the most appropriate thing to do when the detected flux is a function of pixel-phase. I think it probably is (and it might be the only 'proper' thing to do in that event) after talking with John Fowler about it. One thing to keep in the back of your minds is that using a larger support and smoother kernel is not always a bad thing. A little bit of smoothing can make better (and more scientifically viable) images at a small cost to the resolution (... probably should remove this comment if showing my note to Peter!).

Having the MCM built into the code is a nice idea. I think that if you ever want to turn it on for the long wavelength observations, you will need to add regularization. Check out Starck et al. (2002) PASP, 114, 1051 and references for some ideas.

As your data is oversampled, bi-linear interpolation is probably the way to go when you need to interpolate. There is probably a clever way of using a lookup table to speed the process if needed.

I still see a big problem with using weighting by the noise if you derive the Poisson contribution directly from the data. If you estimate the noise from an ensemble (which you might be able to for the CVZ data), then you should be okay. As an example consider two data samples that exactly overlap and have 1 and 3 DN respectively. The noise estimates from each sample are 1 and  $\sqrt{3}$  even though they are both from the same population (whose mean is unknown at this point, let's use three for an example). The coadd is then 1.73 while weighting with the actual sigma ( $\sqrt{3}$ ) gives a value of 2.0. The net result of sigma weighting when you estimate Poisson noise directly from individual data samples is that you systematically bias against the higher measurements.

Frank, your discussion of correlated noise is great. Would you allow us to post a version of it on the SSC website. I think that many Spitzer observers would benefit from it.

## **2.2 *John Good***

My general feeling is that everything looks well in hand. All of the right issues have been identified and are being addressed.

Some potential problems cannot be fully addressed until there is real data in hand, such as what to do about background matching if the flat-fielding is imperfect.

My only real concern is that the amount of computation needed to perform all the processing described may tax the available resources. However, since the processing is intrinsically parallel, at worst this means buying more CPUs.

### 2.3 *Russ Laher*

Based on my review of WISE-project documentation and the material presented at the 11/15/07 peer-review meeting, my assessment is that the current design will meet or exceed subsystem requirements, and the algorithms have been well chosen to perform the required data-processing tasks and have been implemented to provide robustness against non-nominal data.

With regard to selection of a cosmic-ray detection algorithm, which we also discussed at the meeting, I recommend performing completeness and reliability measurements of the candidate algorithms so that the final algorithm can be selected based on a quantitative evaluation. For the Spitzer project, unfortunately, this was not done during the implementation phase, although I did do some of this work later. I found, for example, that there is a drastic tradeoff between C and R for low-coverage cases (e.g.,  $R=0.27$  at  $C=0.9$  and  $R=0.9$  at  $C=0.5$ ). The tradeoff is probably not as severe for high-coverage cases. My recommendation for the WISE project is to at least perform the C&R analysis for a nominal data set and, if time permits, extend the analysis to cover a larger variety of WISE data sets (e.g., various cosmic-ray densities and strengths).

As for design choices of cosmic-ray detection algorithms, using a Laplacian filter at the frame level to enhance cosmic rays will also enhance noise and introduce side-lobe features, which may reduce the performance of multi-frame cosmic-ray detection (assuming the Laplacian-filtered results are fed to the multi-frame outlier detection). On the other hand, local-median background subtraction, with a small window for the median calculation, may be better because the noisy pixels on bright sources will be suppressed by the higher resulting local median value. Ultimately, C&R performance studies will be needed to unequivocally decide which is best.

I mentioned at the meeting that there is a lot of Spitzer software that could be reused by the WISE project. A portion of the software has been ported from the Solaris to the Linux and Mac platforms. †Some additional work may be required to make the software more efficient for WISE images which have 2-4 times as many pixels as Spitzer IRAC images; we are currently doing this sort of work on some of the software for the PTF project (e.g., the flatfield module).

My final comment is with regard to the mottled noise in the coadded image generated by the WISE coadder software. I don't have any experience with the high-resolution algorithm, and so this may be better called a challenge than a suggestion. It seems to me that it might be possible to smooth out the noise in the coadded image by implementing a different windowing scheme, and I am thinking of a method called Taylor weighting, in particular, which has the property of making all side lobes the same height while increasing the main lobe a small amount.

Both you and Frank did a fine job covering the main points at the review meeting, and it is readily apparent that the WISE team has worked diligently to address all project requirements and contingencies. The documents distributed prior to the meeting are up to professional standards, not to mention impressive in their breadth of coverage.

### 2.4 *Dave Shupe*

1. Does the design of the Frame Co-addition program address the requirements on the system?

Yes, with two minor caveats. The first is that the high-level requirements seem to me to be levied on the extractions and not the co-added maps, and there I think there's a little fuzziness as to what are the requirements on the images. The second is that there is no requirement on deconvolution or image improvement, yet the design has been expanded to accommodate the possibility. I'll expand on these two points in more detailed comments below.

2. Are the Frame Co-addition algorithms suitable and appropriate to carry out the system functions?

Yes.

3. Is the design robust to circumstances that will be encountered with the WISE data?

Yes.

Overall, I am impressed with all the hard work and careful thought that was apparent from the review. It is clear that Frank has carefully considered the options for image coaddition, including the previous work that has been done at IPAC for this kind of processing. The AWAIC system is in good hands.

I do have two minor issues (the caveats mentioned above). At the review we spent quite some time on outlier rejection, which confirms how complex this area is. I did not get a good sense that the requirements on outlier rejection at the image level, or on image quality, were defined as clearly as I would like. The higher-level requirements seem to me to apply to the extractions and not to the images. Roc mentioned that image artifacts will be mapped by the IRSA server onto the images (presumably the same as 2MASS). Additionally a multi-frame point-source extractor was mentioned, which in principle should be robust against some image artifacts. So the images don't need to be "perfect" and there is a risk that too many resources will be devoted to outlier rejection.

The second issue is that the high-res algorithm has been baselined for coaddition. As Frank noted, with no iterations the algorithm just produces a "vanilla" coadded image. This is perfectly fine and addresses the requirements appropriately. Since there are no requirements on deconvolution, this is another possible risk area where too many resources could be spent.

I confess that my personal experience with Spitzer has colored my view of deconvolution. Before launch there were a couple of efforts that I knew about for deconvolution of MIPS data, one being the SWIRE team's testing of the proprietary Pixon method, and the other being the development of High-Res by JPL's "Long-Wavelength Center for Excellence" (which I believe is no more). These efforts did not amount to much as far as I know, certainly not the Pixons anyway. However, it may be that the less-than-hoped-for sensitivity of MIPS is the real reason

why there hasn't been much deconvolution activity for Spitzer. A colleague of mine has gotten very excited about his experiments in deconvolving MIPS 70 micron data, so maybe we've neglected this area.

I do make the specific recommendation that Frank check into whatever became of JPL's High-Res effort, and if there is software, to try it out on real MIPS data.

Some other small comments:

- \* Simulations are nice, but if at all possible the AWAIC algorithms should be tried on real data, Spitzer data in particular.

- \* As I mentioned at the review, we were surprised on SWIRE that a few asteroids made it through our outlier rejection - be sure to remember these kind of outliers in addition to the usual rad hits.

- \* I am pleased to see the SIP distortion representation has been baselined for WISE!

- \* As I mentioned at the review, on SWIRE we've regretted making our 24-micron pixels too small (1.2 arcsec mosaic pixels for a 6 arcsec beam, when native detector pixels are 2.45 arcsec and are sampled well enough). The noise is correlated across mosaic pixels which causes headaches for source extractors, and the images become unwieldy. We didn't do investigations to find the "sweet spot" of optimum pixel size, but our feeling is that 2.45 arcsec is near-optimal.

- \* For background matching, I recommend that WISE look carefully at the Montage method.

- \* The goal for background matching is to preserve natural background variations. If zodiacal background has a gradient, does this mean you want to preserve it? I would argue that you don't. Gradients in images make them hard to view. For SWIRE-IRAC, we've removed the gradient while just leaving the mean zody level.

- \* Herschel-PACS and SPIRE are baselining a map-making algorithm known as "MADMAP". I don't know a lot about it, but I think it would be worth it for Frank to chat with some local Herschel people about it. Kevin Xu is the local SPIRE mapping expert, and you could talk to Babar Ali or Dave Frayer about PACS mapping.

## **2.5 *Miscellaneous Comments Recorded During Review***

*Sean Carey* - Are sup-pixel response variations expected for the WISE detectors? If so, what will their effect be on the PSFs?

Answer (from Ned Wright via telecon): It's unlikely. The telescope optical PSF is sampled better than critical by the detector pixels in all arrays and response variations are expected to be negligible.

*Sean Carey* - Consider providing low-resolution coadd products to QA system to help diagnose large scale (i.e. focal-plan scale) problems.

Answer: usefulness will be considered.

*Sean Carey* - Would we consider using a pseudo-median to speed up computations?

Answer: yes, if it's found to be robust and 'accurate' enough.

*Peter Eisenhardt* – How large will the sphere of influence of a bad/masked pixel grow due to the interpolation?

Answer: it's of order the size of the input PRF used for the interpolation. A single masked pixel will not lead to a loss of coverage equal to 1 at the location of the masked pixel, instead, there is an effective loss of "fractional" coverage over the region spanned by the PRF.

*Peter Eisenhardt* – Will rejected pixels be noted in coverage statistics?

Answer: Yes.

*Peter Eisenhardt* – Will PSF kernel interpolation be used when resampling for outlier rejection?

Answer: No. A PSF kernel optimized for accentuating and localizing outliers (e.g., a top-hat) will be used.

*Peter Eisenhardt* – How will statistical accuracy of coadd process be tested? Monte Carlo simulations using simulated sources injected into real data?

Answer: at the poles, where our nominal coverage is 1000 frames, we can create independent co-adds from frame subsets, e.g., 10 co-adds each of 100 frames. We can then assess/check the statistical accuracy, and the computed uncertainties by examining the degree of repeatability across the sub-coadds.

*Nick Gautier* – Will background matching preserve true (astrophysical) structure in the background?

Answer: That's usually the goal, but it depends on what scale. Small scale structure/variability within the proposed scale of an Atlas Image (<~1.5 degrees) are likely to be preserved, but on larger scales (across Atlas Images) it will depend on your background matching algorithm and whether a prior global constraint (e.g., absolute background) is assumed.

*Nick Gautier* – Make sure that we don't do anything in the coadding process that would prohibit matching backgrounds on a global scale, or force the coadd process to be repeated.

Answer: We won't be going back to the frames to match their backgrounds over large (> proposed Atlas Image) scales, and then re-coadding. Matching will first be done locally in order to get each Atlas Image tile as self-consistent as possible for scientific purposes, then between the Atlas Image tiles on larger scales if needed (see below for more details). The WISE image/pixel server will handle the latter.



*John Good* – Will any attempt be made to match the background levels of all of the coadds on a global scale?

Answer: Yes. We will examine the final Atlas Images over the sky and consider whether matching on a global scale is needed. This will only be a problem if a user requests a cut-out (from the pixel server) that falls on the boundaries of Atlas Images. The backgrounds between Atlas Images may not be perfectly consistent. Absolute knowledge of the background (or some other constraint) may need to be folded into the algorithm to avoid non-sense solutions. Also, additional throughput re-scaling will be needed if the photometric zero points between adjacent Atlas Images differ significantly. This is so the user will have the same photometric zero point for all pixels in the queried region.

*Russ Laher* – Does AWAIC accept bit masks as input?

Answer: Yes.

*Jason LaPointe* – Can outlier rejection results be used to test for image compression artifacts?

Answer: Yes indeed.

*Dave Shupe* – Standard deviation images have proven very useful in analysis of Spitzer data.

Answer: Yes. A similar product will be generated as part of the outlier rejection. It will be based on a more robust (outlier resistant) measure than the standard deviation. It is called the MAD (Median Absolute Deviation) image and is rescaled such that it reproduces the sigma of a normal distribution. This image will be used to assess our error model uncertainties, and replace/re-scale them if needed. Therefore, since the uncertainty Atlas Images will be made consistent with the variance from repeated frame observations, there's no longer a need for a separate product. This information will be effectively contained in the uncertainty Atlas Images, and these will be archived.

### **3 SPECIFIC RESPONSES**

#### **3.1 Sean Carey**

*For outlier rejection, using the temporal rejection as Spitzer does should work more than adequately as you have plenty of samples (minimum of 8). Using mad to calculate the sigma is definitely the way to go if the outliers occur on one side of the distribution more than the other (radhits will be positive). One cautionary note is not to regrid the data to too fine a mesh when stacking the input pixels as the interpolation can de-weight the significance of outliers. This may influence your implementation as Frank has proposed a one pass outlier-mosaic scheme*

*which is very I/O efficient. However, if your output grid is too fine, outliers may leak in more than you would like. Gridding the data on at lower resolution for the outlier detection might be a better option(you'll need some data to try it on).*

Response:

The MAD measure is insensitive to the location of outliers, i.e., either left or right of the mode of a distribution. Remember, the MAD represents the Median of the Absolute Deviations of all samples from the median. The fact that absolute differences are used ensures that both sides of the distribution are treated equally. However, the median can be slightly biased towards the tail that has the most outliers (in extreme circumstances). This causes the MAD measure and hence sigma computed there from to be biased slightly high. Given >8 samples, this bias will be negligible, but nonetheless, it will be assessed using an iterative sigma clipping method.

*In any event, you may want to consider "growing" the outlier masks to catch the edges of the radhits. For IRAC, the Si:As arrays tend to have fatter radhits (as the detectors are physically thicker). For these arrays, I used a low fraction of radhit overlap to output pixel overlap as a threshold. It's conservative, but that's okay as you have plenty of samples. While I am thinking about it, you will not want to apply the image construction kernel weighting for the outlier detection gridding and use straight area overlap instead.*

Response:

Yes. Every pixel will be considered for "outlier status" and "fat radhits". I can't see how these will be missed using the generic method I proposed. Also, I do not plan to use the image construction kernel for the outlier detection gridding. It will be one optimized for detecting outliers, most likely a top-hat PRF (= area weighting).

*Most asteroids will probably only be partially rejected as you won't use a low enough sigma threshold to reject the wings of the PSF for the asteroids. That's okay, but you might want to have a flag in the QA or post-processing stage to indicate that this is happening. The cores of the asteroids should be well rejected though as they will be significant positive outliers in 50% or less of the coverage. That's what we noted in MIPS GAL in any event.*

Response:

Very good point. Moving objects will only be (totally) flagged if they are moving at speeds  $> \sim \text{FWHM}(\text{PSF})$  per nominal 94 minute orbit. Objects moving slower than this (or appreciably so) will survive on the co-add. As you said, these will only be partially rejected. In the end, these "moving remnants" will appear elongated in a co-add, i.e., if they move at a speed such that their PSFs overlap in a line. But we're not concerned about this. These (supposedly fainter) asteroids may get included in the source detection list and all for the better - this will allow us to measure their fluxes using PSF fitting off the frames. Owing to their motion, the overall  $\chi^2$  across frames will be larger, and this may then serve as a flag for QA. We don't intend to surgically eradicate all instances of moving objects from the co-adds.

*Background matching for large maps is a challenge. One thing you will have to decide upon is whether you are going to remove a model of the Zodiacal light or not. If you are going to provide a product that could be used for extended structure studies, I would recommend considering zody removal (especially if you get the second epoch data!).*

Response:

There is no plan to support extended structure studies on large scales, i.e.,  $> \sim 1$  degree (typically the size of an Atlas Image tile).

*I'm attaching a draft of the MIPS GAL 24  $\mu$ m processing paper by Don Mizuno. Section 5 has an extensive discussion of our application of the overlap correction. You could implement a similar solution using the sparse matrix direct solution that Don did or a relaxation method as I think Montage does. The important thing in the MIPS GAL background solution is the damping terms applied. Please do not distribute this paper as we have not yet submitted it.*

Response:

Very nice method. I agree that without a global constraint (provided by the actual data or otherwise), your solutions would veer into the blue.

*Having the MCM built into the code is a nice idea. I think that if you ever want to turn it on for the long wavelength observations, you will need to add regularization. Check out Starck et al. (2002) PASP, 114, 1051 and references for some ideas.*

Response:

Thanks. I want to remind you that deconvolution (HIRES'ing) is not in the baseline plan. My intent is to enhance this functionality for offline use (as a research topic) after the dust has settled with the primary WISE products.

*I still see a big problem with using weighting by the noise if you derive the Poisson contribution directly from the data. If you estimate the noise from an ensemble (which you might be able to for the CVZ data), then you should be okay. As an example consider two data samples that exactly overlap and have 1 and 3 DN respectively. The noise estimates from each sample are 1 and  $\sqrt{3}$  even though they are both from the same population (whose mean is unknown at this point, let's use three for an example). The coadd is then 1.73 while weighting with the actual sigma ( $\sqrt{3}$ ) gives a value of 2.0. The net result of sigma weighting when you estimate Poisson noise directly from individual data samples is that you systematically bias against the higher measurements.*

Response:

Yes - all in agreement. The assumption of using  $\sqrt{\text{counts}}$  as an estimate of the standard deviation of the ensemble from which that realization was drawn can be hugely biased. I'm surprised observers still use this (for weighting purposes), even when many observations of the same scene are available. However, it's the 'best' you can do when few observations are

available. It makes more sense to use the  $\sqrt{\text{mean}(\text{counts})}$  where  $\text{mean}(\text{counts})$  is the mean of the ensemble

(repeated observations). But, this latter estimate only tells you about the random (Poisson) component, not systematics. My plan is to use the variance in the repeated (outlier corrected) pixel samples as a measure of the uncertainty where ever possible, and rescale the error model uncertainties using  $\chi^2$  tests before using them downstream. These final (rescaled) uncertainties are what will be archived along with the frame data. In cases where there are few samples ( $\ll 6$ ) in a stack, I will look at how the  $\chi^2$  of the neighboring pixels fair before rescaling. This functionality will be part of the outlier rejection and my apologies if I didn't make it clear at the review.

*Frank, your discussion of correlated noise is great. Would you allow us to post a version of it on the SSC website. I think that many Spitzer observers would benefit from it.*

Response:

I have no problem you posting it. There is a new version of the SDS at: <http://web.ipac.caltech.edu/staff/fmasci/home/wise/sds-wsdc-D005-v2.8-awaic.pdf>

### 3.2 Russ Laher

*With regard to selection of a cosmic-ray detection algorithm, which we also discussed at the meeting, I recommend performing completeness and reliability measurements of the candidate algorithms so that the final algorithm can be selected based on a quantitative evaluation. For the Spitzer project, unfortunately, this was not done during the implementation phase, although I did do some of this work later. I found, for example, that there is a drastic tradeoff between C and R for low-coverage cases (e.g.,  $R=0.27$  at  $C=0.9$  and  $R=0.9$  at  $C=0.5$ ). The tradeoff is probably not as severe for high-coverage cases. My recommendation for the WISE project is to at least perform the C&R analysis for a nominal data set and, if time permits, extend the analysis to cover a larger variety of WISE data sets (e.g., various cosmic-ray densities and strengths).*

Response:

Simulations to assess C&R and optimize thresholds/outlier windows will be performed during the implementation phase. An assessment against other outlier detectors (e.g., MOPEX) on real data will also be made.

*As for design choices of cosmic-ray detection algorithms, using a Laplacian filter at the frame level to enhance cosmic rays will also enhance noise and introduce side-lobe features, which may reduce the performance of multi-frame cosmic-ray detection (assuming the Laplacian-filtered results are fed to the multi-frame outlier detection). On the other hand, local-median background subtraction, with a small window for the median calculation, may be better because the noisy pixels on bright sources will be suppressed by the higher resulting local median value. Ultimately, C&R performance studies will be needed to unequivocally decide which is best.*

Response:

There was some misunderstanding here. I was planning (and it's still TBD) of using a gradient filter to first detect/flag outlier spikes in the frames themselves. This is a rudimentary first pass, and is decoupled from the second more rigorous pass that uses the frame stacks.

*I mentioned at the meeting that there is a lot of Spitzer software that could be reused by the WISE project. A portion of the software has been ported from the Solaris to the Linux and Mac platforms. †Some additional work may be required to make the software more efficient for WISE images which have 2-4 times as many pixels as Spitzer IRAC images; we are currently doing this sort of work on some of the software for the PTF project (e.g., the flatfield module).*

Response:

I am aware of all the software written for Spitzer. We will strive to use or upgrade any IPAC/SSC heritage software that's fit for our purpose.

*My final comment is with regard to the mottled noise in the coadded image generated by the WISE coadder software. I don't have any experience with the high-resolution algorithm, and so this may be better called a challenge than a suggestion. It seems to me that it might be possible to smooth out the noise in the coadded image by implementing a different windowing scheme, and I am thinking of a method called Taylor weighting, in particular, which has the property of making all side lobes the same height while increasing the main lobe a small amount.*

Response:

Smoothing the co-add images will certainly increase their aesthetic appeal, but this would hide the underlying noise structure and make it difficult to quantify uncertainties (and variations there-of) over the co-add.

### **3.3 Dave Shupe**

1. Does the design of the Frame Co-addition program address the requirements on the system?

*Yes, with two minor caveats. The first is that the high-level requirements seem to me to be levied on the extractions and not the co-added maps, and there I think there's a little fuzziness as to what are the requirements on the images. The second is that there is no requirement on deconvolution or image improvement, yet the design has been expanded to accommodate the possibility. I'll expand on these two points in more detailed comments below.*

Response:

Right, requirements on the co-add images (and ancillary products) only pertain to their content, format, and delivery methods. However, it is really driven (in an indirect way) from the requirements on source detection. The ability to meet/surpass these will depend on co-add image quality, and in turn, on all other instrumental/artifact mitigation steps upstream. The deconvolution (HIRES) algorithm was added to support future research. However, this is not in the baseline plan, and no further enhancements on this aspect of the code will be made until well after the end of on-orbit operations.