The NEWFIRM Quick Reduce Pipeline and Data Analysis Tools

Last update: 13 October 2010

These notes provide a quick guide to the NEWFIRM Quick Reduce Pipeline and other data reduction software and tools for use at the telescope. NEWFIRM is currently operating at the Blanco 4m telescope at Cerro Tololo, and these notes have been updated for operations at CTIO.

Table of contents:

Quick reference information:

- ### Contacting the NOAO NEWFIRM pipeline group

  The NOAO Data Products NEWFIRM pipeline group is Rob Swaters (University of Maryland), Frank Valdes (NOAO), and Mark Dickinson (NOAO). To contact us, please send e-mail to nfpipe@noao.edu.

- ### The Pipeline Review Log

  Once you have started observing, the first place to look for the results of pipeline processing is the Pipeline Review Log. This, and other ways to examine data products from the NEWFIRM Quick Reduce Pipeline, are described in more detail below.

The NEWFIRM Quick Reduce Pipeline (QRP) - an introduction

The NEWFIRM Quick Reduce Pipeline (QRP) operates on computers at the CTIO 4m telescope, and provides basic, automated data processing for NEWFIRM during the course of an observing run. It is designed to provide basic reduced image products and data quality feedback to the observer on a reasonably short timescale as the observations are being taken. The QRP is not intended to produce science-quality reductions. It takes several shortcuts to speed data processing, such as using only single-pass sky subtraction. Moreover, the best possible calibration data may not be available at the time a given observation is taken. For example, frequently an observer will not take darks with exposure times appropriate for their science observations until after the science data are taken. In this case, the pipeline will fall back on earlier darks from a library, using something with the closest available exposure time.

The QRP is triggered at the end of each NEWFIRM exposure. However, in general NEWFIRM observations are taken in sequences of multiple exposures that are defined with the NEWFIRM Observing Control System, or NOCS. Examples include sequences of darks or flats, or dithered sequences of scientific observations. The QRP will not process observations taken with 'TEST' sequences (as created by the NEWFIRM script-generating GUI), or focus sequences (see the discussion of the nffocus routine below for information on analyzing focus data).

Much of the QRP processing does not begin does not begin until the last exposure in an observing sequence is completed. For night-time science observations, the pipeline applies dark and flat calibrations, sky subtraction (usually but not always based on a running median), determines a WCS and photometric calibration relative to 2MASS stars in the field, resamples images to a common pixel grid, and combines them into a stacked mosaic. The QRP will also process and combine darks and dome flat sequences.

Processing of the data starts immediately upon completion of a sequence. For many observing programs, the typical cadence of exposures during night-time science observations will be of order one image per minute. At this data rate, the pipeline will typically take somewhat longer to process the data than the time needed to complete the observing sequence itself. However, if you are observing with a much quicker cadence (e.g., a long series of 15 second exposures), the pipeline may not be able to keep up.

The products of the Quick Reduce Pipeline are *not* archived or saved in any other way. If you want to keep the QRP-produced FITS products for your own use, this is your responsibility. The pipeline review web pages that you can examine during your run will also be removed at the start of the next observer's run. Note also that these review pages are also behind the mountain firewall, and are thus not accessible to people outside the mountain internet domain.

A more detailed description of the QRP and list of some known caveats is given below.

How to observe in a "pipeline friendly" manner

There are a few important steps that you can take to ensure that your observations are done in a way that ensures good performance from the Quick Reduce Pipeline.

Take at least *some* calibrations before you start to observe: The QRP operates in pseudo-real time, processing your observations after the completion of each observing sequence. Therefore, it needs to have appropriate calibration images - specifically, darks and dome flats - available at the time the observations are taken. The QRP maintains a calibration library, and will look back in that library for the most appropriate calibrations taken as recently as possible. However, if they are not available (e.g., flats for the particular filter you are using), then it will simply skip this processing step. Or, it may use rather old and inappropriate calibration files (e.g., darks with the wrong exposure times) as its "best guess". The best thing you can do is to take some calibrations *before* you observe, if at all possible. Note that on the second (and subsequent) nights of your run, the QRP can fall back on calibrations from the first night.

- **Dome flats:** At present, the QRP *only* works with dome flats, not sky flats. You *must* take dome flats both with the lamps on and with the lamps off in order for the QRP to process them correctly (it combines each group, then takes the difference between them). *To ensure that this is done correctly, you should use the **DomeFlatSequence** (or DFLATS) option from the NGUI observing script generator, and follow the instructions carefully about when to turn on and off the dome flat lamps.* We *strongly* discourage the use of the *DomeFlat* sequence instead, because this does not provide reliable information to the pipeline about the flatfield lamp status, and it complicates the issue of grouping the "lights on" and "lights off" flats. Note also that you must take *at least three exposures* in each of the lights-on and lights-off status in order for them to be processed by the QRP.

  Ideally, you should also take darks with exposure times appropriate for your dome flats. However, this is not critically important, especially for the QRP, since the dark signal is essentially removed by the lights on - lights off differencing.

- **Darks:** Good IR observing practice requires that you take darks with the same combinations of exposure time, number of coadds, number of digital averages, and number of Fowler samples as you use for each type of observation that you will take at night, including standard stars, etc. Unfortunately, it is often difficult to know exactly what exposure time you will use before the night begins, due to the variable nature of the near-infrared sky background. However, we encourage you to take a good guess at your exposure times, coadds, etc., and take *at least three* dark exposures before you start the night. (The QRP will not combine sequences of fewer than three darks or flats.) If the QRP cannot find darks with the exact exposure times (or other parameters) needed, it will use whatever else it can find in its calibration library with parameters as similar as possible. However, given the variety of parameters at play (including exposure time, digital averages, etc. etc.), this matching is a very inexact procedure.

Note that the NEWFIRM science pipeline, which will process your data after your observing run is over, is not restricted to using calibrations taken *before* your science observations. It will comfortably use darks and flats taken after your night of observing, or on the next day of your run, etc. However, it is always good infrared observing practice to take flats and darks as close in time as possible to when you make your science observations.

Dither, don't linger: The NOCS observing sequences that you create with the NGUI script generator will in principle allow you "linger" at a given position in an observing sequence, taking more than one exposure per dither position, by setting "NumObs" to a value greater than 1. At present, however, the QRP does not process such data correctly, and in particular the sky subtraction will be badly corrupted. We will change this in the future, but in the meanwhile (and, in fact, in general) we would recommend that instead you use Coadds (set "coadds" to >1) instead of "NumObs". This will take several exposures and coadd them on board the instrument before writing to disk. This will also improve your observing efficiency - what's not to like?

Keep dithering: The pipeline uses a running median sky subtraction procedure as long as a minimum of *4 dithered exposures* are taken in a sequence. For shorter sequences, it will use simple pairwise image subtraction, and you will see pronounced negative residuals in your images. From the point of view of sky subtraction, longer sequences with more dither positions are better - more than 9 if possible, but the longer, the better. If your science requires shorter dither sequences, then by all means go ahead and use them. However, if you plan to take many dithered exposures on a given field, it is better to do so in longer observing sequences rather than to break them up into many short sequences of a few frames each.

Offset sky fields: If you are observing a large, extended object that fills a significant fraction of the NEWFIRM field of view, or perhaps if you are observing a very crowded field, then you may wish to "chop" (perhaps frequently, perhaps occasionally) to a relatively blank field that is offset a significant distance from your target. The pipeline can process such data, constructing its background images using only the frames taken pointing at the offset field. However, for this to work, the pipeline needs to know which frames are the "object" exposures and which ones point at the offset "sky" field. This is *only* possible if you observe using the appropriate type of sequence generated with the NGUI tool. The offsets to blank sky are done using the RA and Dec Offset distances in the Telescope Configuration section of the NGUI. There are a variety of script options that will observe in this manner, including DeepRich, ModMapRich, and QuickMapWithSky. You should consult with Ron Probst or someone else familiar with the instrument to discuss which is best for you. Any of these sequences, however, will record metadata in appropriate header keywords which will tell the pipeline which frames are "object" and which are "sky". Note that it is possible, but not advisable, to create your own customized offsetting procedure using Map offsets in other NOCS/NGUI scripts (e.g., DeepSparse or QuickMap), rather than the Telescope Configuration offsets. However, the pipeline will *not* correctly process these data.

One special case is the "4-shooter" or "4Q" observing mode, which can be invoked by selecting the "4Q" dither or map option within most scripts. This is designed for observations of an extended object which fills a good portion of one NEWFIRM quadrant but not the whole 4-detector NEWFIRM field of view. 4Q mode places the target near the center of one quadrant, and then cycles the target through each of the four quadrants as many times as you wish, optionally applying other (usually smaller) dither offsets so that the target positions do not exactly repeat from cycle to cycle. 4Q-mode correctly records in the image headers the information about which detector includes the target, and the NEWFIRM pipeline will correctly process these data, constructing sky frames for each detectors using only those exposures which do *not* include the target. This has proven to be a very effective and efficient way to observe extended objects that are still smaller than 14 arcmin or so in diameter (i.e., the field of view of a single NEWFIRM detector).

Test exposures: If you want to take a short test exposure for some reason (e.g. to check instrment aliveness, telescope pointing on a bright star, check the sky brightness, etc.), use a *Test* sequence from the NGUI script generator. Observations taken with the *Test* sequence will not get processed -- the QRP will not do a good job processing single test exposures anyway.

Observer setup information

NEWFIRM observers generally use the screens for the computer ctiozm to operate the instrument and examine incoming data. These screens get very crowded, so we recommend using another computer (or at least a different desktop on *ctiozm*) for examining pipeline output. There is another computer in the 4m control room, *ctioa8*, which would be suitable, or you could work from a laptop.

On the ctiozm console (data-taking computer):

- On the left-hand *ctiozm* screen, where you are connected to the data handling computers via VNC, go to the DHS Supervisor window and select the "Paths and Files" tab. There, you can give the path to the directory where your raw data will be written, as well as a path to the directory where the pipeline will deliver reduced FITS images. You might wish set the latter to be a subdirectory (e.g., called "QRP") of your data directory for each night, or, you could create one directory for the QRP products from your whole run. Be sure to hit *<enter>* after typing the paths in this window. From an xterm or shell window, create the directory that you specified for the QRP products.

On another computer (e.g., ctioa8, or from a laptop):

- Open a web browser and point it to the NEWFIRM pipeline review log. This is a web page showing a log of data that the QRP has processed for your run so far. Inititally, it will be empty until you take data for the QRP to process, such as dome flats, darks, or your first night-time observing sequences (other than test or focus observations). The QRP review web pages are discussed in more detail below.

- Open one xterm window and log in to the *newfirm* computer via *ssh observer@newfirm-ct*. (You should have received the password from the instrument scientist who got you set up at the start of your run.) From there, you can run the plstatus command to check that the pipeline is running properly.
- Open another xterm window and log in to the *nfpipe-01* computer via *ssh observer@nfpipe-01-ct-ct*, using the same password as for *observe@newfirm-ct*. From there, you can run the dqquery command, which provides data quality information for images that the pipeline has processed recently.
- Open an xgterm window and log in to the *newfirm* computer via *ssh -X observer@newfirm-ct*. Open an image display (ximtool or ds9) in the background, and then start an IRAF session from the /home/observer/ directory by typing 'cl'. Change directories to go to the location you specified for the QRP FITS output products. You can then display and otherwise analyze these images as they appear during the course of the night.

Checking the Pipeline status

You can use the command **plstatus** To check the operational status of the NEWFIRM Quick Reduce Pipeline from an xterm window logged in as *observer@newfirm*. If the pipeline is functioning properly, you should expect to see this:

```
[observer@newfirm ~]$ plstatus

Pipeline node status report: Sat Mar 15 15:24:44 2008
Current node is nfpipe-01

Node         Connected    NodeMgr      Available Pipelines
-----------------------------------------------------------------
nfpipe-01    Yes          Running      NEWFIRM:
                                          wcs, mtd, gcl, flt, drk,
                                          fil, dts, ndp, gos, sfl,
                                          mkd, swc, sky, sgc, stk,
                                          sdk
-----------------------------------------------------------------
nfpipe-02    Yes          Running      NEWFIRM:
                                          wcs, drk, flt, fil, gos,
                                          sfl, sky, sgc, sdk, swc
-----------------------------------------------------------------
```

If, instead, the machines are up but the pipeline is not running, you should see this:

```
[observer@newfirm ~]$ plstatus

Pipeline node status report: Sat Mar 15 15:24:44 2008
Current node is nfpipe-01
```

```
Node         Connected    NodeMgr      Available Pipelines
------------------------------------------------------------------
nfpipe-01    Yes          No
------------------------------------------------------------------
nfpipe-02    Yes          No
------------------------------------------------------------------
```

If "Connected" reads "No" then the pipeline machine itself is not up.

Examining the pipeline data products

Before observing, you should use the DHS GUI you should specify the path for a directory where you would like the QRP FITS data products to be delivered (see setup instructions above). We recommend that you choose a directory other than that where your individual raw images will be written. For example, using a "QRP" subdirectory would be appropriate.

You can inspect the results of the QRP processing in two ways: via the data product review web pages that are generated by the pipeline, or by direct examination of the FITS data products with your favorite image processing tools (e.g. ximtool, ds9, IRAF, IDL, etc.).

Pipeline review web pages:
The QRP produces web pages that provide easy access to pictures of your reduced data, as well as various diagnostic information generated by the pipeline. You can find these QRP review web pages at the following URLs:

- For browsers anywhere within the NOAO domain: http://nfpipe-01-ct.ctio.noao.edu/NEWFIRM_pipeline/observer/
- Full IP address in case the DNS fails to resolve the name: http://139.229.14.130/NEWFIRM_pipeline/observer/

Note that you will need to reload/refresh this web page in your browser periodically, as it will be updated as new data are processed and added to the QRP review log.

You will see a log of data that the QRP has processed for your run so far, with various information about each observing sequence, including:

- image filenames
- UT at the start of the observing sequence
- observation type
- exposure times
- filter
- focus setting
- number of digital averages (DA) and Fowler samples (FS)
- airmass

- image title

In addition, this log reports the ranges spanned by some data quality parameters that are measured by the pipeline, such as:

- sky brightness (in ADU and mag per square arcsec)
- seeing FWHM
- magnitude zeropoint estimated from 2MASS stars (useful for monitoring relative transparency)
- estimated 5-sigma point source sensitivity

Clicking on the image filenames for a given sequence takes you to a pipeline review page with information about data from that sequence. There are review pages for processed, combined dark frames and dome flats, as well as for night-time on-sky observations. For night-time science observations, the QRP review pages provide:

- thumbnail pictures of the data
- links to images at two different higher resolutions,
- basic header information about each image
- links to the full FITS headers
- data quality measurements (seeing, magnitude zeropoint, depth, etc. - see above)

For dithered on-sky observing sequences, preview images are shown for each individual exposure, as well as for a stacked mosaics of the observations. The stacked images are shown in the second column of thumbnail images. If, during the course of an observing sequence, the telescope motions exceed a particular hardwired tolerance (currently 8 arcmin, but this may change), the QRP will break the sequence into multiple, independent stacks that will all appear in the 2nd column of thumbnail images.

FITS data products:
When processing is complete for a sequence, FITS files of the combined image (the "mosaics" or "stacks") will be delivered to the directory that you specified (see above). The names of the files start with NFQR (for NEWFIRM Quick Reduce). FITS files are delivered only for on-sky observations, not for dome flats or dark calibrations.

DQQUERY - monitor seeing, transparency and sky brightness

dqquery is a command that displays measurements of the seeing, transparency, and sky brightness made by the NEWFIRM Quick Reduce Pipeline. The pipeline creates a catalog of sources from which the average sky brightness (ADU) and seeing (arcsec) are determined. The catalog is matched against 2MASS sources and a magnitude zero point is estimated. Naturally, if the exposure contains few or no stars then the seeing and zero points may be indeterminate or poorly determined. The matching to the 2MASS catalog may also fail for similar reasons or due to large errors in the telescope coordinates.

The measurements are recorded in a pipeline database which is queried by **dqquery**. The results are written when the command is first executed and then every 5 minutes thereafter until the command is terminated. A control-C termination will print a couple of lines of errors which can be ignored. The fact that a database query is made means that it is possible for the database to be unavailable during periods when the pipeline is busy recording data. So just wait or try again.

To use **dqquery**, open an xterm window on the *newfirm* computer (e.g., via *ssh observer@newfirm-ct*). There, execute the command **dqquery**. The output is illustrated below, along with an example of the program output. The listing gives a number of recent measurements with most recent first. Depending on the filters used there may be information from different filters. There are a few user selectable parameters, particularly the number of recent measurements desired.

DQQUERY example:

```
[observer@newfirm ~]$ dqquery -h
Usage: /shared/pipeline/V1.0/MarioSrc/bin/dqquery [OPTIONS]
Display the seeing, photometric zeropoint, and background levels
for recent exposures.
    -a, --average      NOT IMPLEMENTED YET
        Average data quality parameters over the requested time interval.
    -d, --debug
    -f, --filter
        Only display data quality parameters for the requested filter.
    -g, --graph                       NOT IMPLEMENTED YET
        Display the data quality parameters in graphical form.
    -h, --help
    -n, --noloop
        Execute query only once, do not enter endless loop
    -r, --recent
        Show the data quality parameters for the requested number of
        exposures. The default value is 16.

[observer@newfirm ~]$ dqquery

-------------------------------------------------------------------------
-------
        Name        Date/Time  Seeing   Mzero      Bkg  Filter
MJD
                               arcsec     mag      ADU
-------------------------------------------------------------------------
-------
   image4651  20071121T120953    0.85  26.202   3755.9      J
54425.5068634
   image4650  20071121T120912    1.06  25.814   1843.2      J
54425.5063889
   image4649  20071121T120831    0.90  25.775   1835.9      J
54425.5059143
   image4648  20071121T120755    0.97  25.795   1828.6      J
54425.5054977
   image4647  20071121T120715    0.89  25.786   1839.0      J
54425.5050347
```

```
   image4646  20071121T120639    0.90   25.733   1842.9       J
54425.5046181
   image4645  20071121T120558    1.00   25.893   1873.3       J
54425.5041435
   image4644  20071121T120522    0.88   25.835   1853.2       J
54425.5037269
   image4643  20071121T120441    0.77   25.811   1870.9       J
54425.5032523
   image4642  20071121T120406    0.80   25.810   1827.3       J
54425.5028472
   image4640  20071121T120111    0.94   25.847   1895.7       J
54425.5008218
   image4639  20071121T120030    0.86   25.768   1900.2       J
54425.5003472
   image4638  20071121T115954    0.96   25.724   1868.2       J
54425.4999306
   image4637  20071121T115914    0.92   25.845   1879.5       J
54425.4994676
   image4636  20071121T115838    0.98   25.800   1856.3       J
54425.4990509
   image4635  20071121T115758    0.94   25.753   1874.5       J
54425.4985880
-------------------------------------------------------------------------
-------
```

One useful application of **dqquery** is to dump a log of data quality measurements for all observations from a long stretch of time (e.g., for the last night). You can do something like this:

```
[observer@newfirm ~]$ dqquery -n -r 500 > dqquery_log.txt
```

The **-n** option turns off the automatic looping display, while **-r 500** dumps the data quality information for the last 500 images that the pipeline has processed. At present, there is no simple way to request DQ information just from a particular night or range of times, but it is easy to just dump measurements from a large number of files and then edit the output to restrict to the desired image names or data/time range.

NFFOCUS - NEWFIRM focus evaluation IRAF task

From within an IRAF session the task **nffocus** may be used to analyze a focus sequence. It has been designed to make measurements and select the best set for estimating a best focus. There is a description of **nffocus** under the IRAF help system:

```
cl> phelp nffocus
```

This task has been designed to work conveniently with the standard NEWFIRM focus recipe. In this recipe a sky offset exposure is first taken followed by a sequence of exposures on the focus field where the focus is changed in uniform steps. After completion of the sequence the task only needs to identify the set of exposures. This can be done by specifying either a single image number or image name within the sequence or by a list of focus sequence images. In the latter case it is convenient to make a file containing the list and specify this to task with the @file syntax.

```
cl> nffocus 17142          # Image number
cl> nffocus obj17142       # Root name is chosen by the observer
cl> nffocus @focusj.list   # Prepared list of images
```

The task will initially make catalogs of the sources from the sky subtracted focus field images. This will take a few minutes. Repeating the task will simply reuse these catalogs so it is quick to return to the graphical analysis. In this graphical analysis you can delete sources, view spatial trends, and visually confirm the automatic best focus recommendation. There are many options but the initial FWHM verses focus plot is often all you need along with the 'd' key to delete points to clean up the outliers. The '?' will give you a quick summary of the available cursor keys and commands.

Quick Reduce Pipeline: details and caveats

The QRP is designed to quickly process data taken at the telescope to provide feedback to the observer about the progress of their observing program. It also reports various quality measurements to help the observer judge the observing conditions. The QRP intentionally takes some shortcuts in order to ensure faster processing and quick feedback to the observer.

The processing steps that the QRP applies to on-sky science observations are:

- **Dark/bias subtraction:**. If dark frames with the same (or similar) exposure time is available, the QRP will subtract these from the science data. These dark exposures are selected from the QRP's database of calibration data, and may come from a different night or even a different observing run. If no suitable dark is found, no dark is subtracted, but processing will continue. To first order, the dark will be subtracted during the sky subtraction stage.
- **Nonlinearity:** The NEWFIRM arrays are increasingly nonlinear up to about 10000 ADU, where the nonlinearity reaches approximately 5%. The nonlinearity increases very steeply above 10000 ADU, and it may not be possible to accurately calibrate pixels brighter than that level. We have written a draft report on NEWFIRM nonlinearity and the method that the pipeline uses to correct it. The QRP is using the correction described in this report, which should correctly take into account the additional signal accumulated in the interval between pixel reset and first read, which can be signifiacant and which varies over the array.
- **Flat fielding:** If a flat field in the appropriate filter is available, the QRP will divide it into the science exposure. The QRP may use flats from a previous night or even a previous observing run. If no suitable flat is found, processing will continue without flat fielding.

- **Sky subtraction:** For long sequences (5 exposures or more), the sky is subtracted by calculating a running median with a window of up to 9 exposures. A simple but robust form of outlier rejection (similar to the 'pclip' algorithm in IRAF imcombine) is used to minimize the impact of sources on the sky images. For shorter sequences, sky subtraction is done by subtracting the closest image in time but within the same sequence (i.e., pairwise sky subtraction). For sequences of one, the sky is subtracted by fitting a smooth polynomial to the background after pixels well above the mean background level (e.g., due to sources) have been excluded. If offset sky exposures are taken *as part of the sequence*, then these will be used for sky subtraction.
- **Astrometric calibration:** After sky subtraction, the QRP makes a simple catalog of sources detected in the image. A list of nearby 2MASS stars is retrieved based on the nominal telescope pointing information, and the two lists are matched to derive a world coordinate solution (WCS). Nonlinear terms of the geometric distortion solution have been derived from previous astrometric calibration observations and are used in this procedure.
- **Photometric calibration:** Measurements of unsaturated, matched 2MASS stars are used to determine a rough photometric zeropoint (MAGZERO) for each image.
- **Bad pixel masking:** A static bad pixel mask in the calibration library is used to mask known pixel defects before images are combined in a dither stack.
- **Stacking:** The individual, WCS-calibrated images from an observing sequence are reprojected onto a common tangent plane, taking into account the nonlinear geometric distortion solution. Linear interpolation is used when resampling the images, for efficiency. The photometric zeropoint derived from the analysis of 2MASS stars is used to rescale all images to a common intensity scale, and the reprojected images are then combined into single stacked image as long as the dither and mapping offsets are small compared to the NEWFIRM field of view. If the telescope motions during a sequence span a wider extent than a particular tolerance, the QRP will break the sequence into multiple, independent stacks.

The QRP makes several shortcuts with respect to the full processing done by the NEWFIRM science pipeline. We comment on some of these here, and also on some general processing issues that apply to both the QRP and the science pipeline.

- **Coadding, Fowler sampling, etc.:** The QRP does not take full account of on-board image coadding or changes in Fowler sampling. In general the processing for such data should be adequate for purposes at the telescope, as the header exposure time information does take proper account of the coaddition. The science pipeline matches calibration frames (mainly darks) by coaddition and Fowler parameters as well as by exposure time, but this is not done in the QRP.
- **Sky subtraction:** The QRP does not mask sources when constructing sky frames from dithered on-sky data. A simple clipping algorithm is used that should prevent the brighter pixels from sources from affecting the sky image, but fainter sources (and the fainter "wings" of brighter sources) will tend to be imprinted on the sky images, and therefore negative "shadows" of sources from other dither positions will often be evident in a sky-subtracted image. For observing sequences with fewer than 5 images, pairwise differencing is used, and these negative images will be even

more prominent. The science pipeline carries out multiple passes of sky subtraction, masking sources in later stages.

- **Lingering:** Sky subtraction is not configured to handle observations where repeat exposures are taken at a given dither position. In general, such an observing strategy is not recommended: you should dither the telescope between each image, over as many independent dither positions as possible, to ensure good sky subtraction. Internal NEWFIRM coaddition can be used to obtain longer exposure times per dither position, in which case the QRP will handle sky subtraction correctly. However, if for some reason your observing strategy requires you to take multiple exposures (individually written to disk) per dither position, the pipeline will not process the sky subtraction correctly.
- **Persistence:** The NEWFIRM arrays exhibit latent images which can persist for many minutes after a bright source is observed. As one dithers, brighter sources will leave a "trail" of slowly fading latent images behind them. In image stacks, small "constellations" of these persistent images will be visible. The NEWFIRM science pipeline masks persistent images for some period of time after each exposure, but this is not done in the Quick Reduce Pipeline.
- **Nonlinearity and saturation:** The correction of NEWFIRM nonlinearity is described briefly above, and in a draft report. The correction that is currently applied uses a single linearity correction (essentially, a quadratic function) for the whole array, although it takes into account the different readout timing for each pixel (see the report for details). This is likely to be an oversimplification, and we plan to measure linearity coefficients for each pixel and employ these in the future. Perhaps more significant is the treatment of saturation, which turns out to be inherently difficult to detect, especially for short exposure times and high count rates. The QRP is currently oversimplifying this grossly by simply flagging pixels with >10000 ADU as saturated. In practice, the saturation level varies as a function of the exposure time, the number of digital averages and Fowler samples, and with position on the array.
- **Photometric calibration:** For both the QRP and the science pipeline, no color terms are employed when deriving MAGZERO from the analysis of 2MASS stars. For filters other than the standard broad JHKs bands, MAGZERO may not be strictly meaningful. The calibration uses the 2MASS magnitudes from the broad band closest to the central wavelength of the filter in which the observations are taken, and therefore may be reasonably close for filters that are not far from the standard broad bands. In any case, relative variations in MAGZERO can be useful to track photometric transparency during an observing sequence.
- **Photometric and image quality screening:** The QRP measures relative magnitude zeropoints per image (MAGZERO), and as described above, it uses this information to rescale images before combining them (e.g., to take into account variable cloud cover). However, it does not weight images differently depending on the data quality, e.g., transparency, sky brightness, or image quality, nor does it exclude images which might actually degrade the net quality of the combined stack.
- **Stacking:** When the astrometrically calibrated images are stacked, the pixel values are combined as a median with outlier clipping. This is a quick and easy way to reject outliers due to unmasked pixel defects, cosmic rays and other problems, but it

may not be optimal or appropriate in conditions where seeing or transparency are varying significantly.

Cadence: The pipeline may have a hard time keeping up with observations with a very high cadence, i.e., long series of short exposures. Images in filters with high background count rates (i.e., the H and Ks broad bands) require short exposure times to avoid saturation. You should use internal coadding to take several exposures before writing them to disk, dithering perhaps every minute or more. This will increase observing efficiency and facilitate the pipeline operations.