Publications of the Astronomical Society of the Pacific, 129:105003 (13pp), 2017 October https://doi.org/10.1088/1538-3873/aa83fd
© 2017. The Astronomical Society of the Pacific. All rights reserved. Printed in the U.S.A.

OPEN ACCESS

Improved Reference Sampling and Subtraction: A Technique for Reducing the Read Noise of Near-infrared Detector Systems

Bernard J. Rauscher[1], Richard G. Arendt[2], D. J. Fixsen[3], Matthew A. Greenhouse[1], Matthew Lander[7], Don Lindler[4], Markus Loose[5], S. H. Moseley[1], D. Brent Mott[6], Yiting Wen[6], Donna V. Wilson[7], and Christos Xenophontos[7]
[1] NASA/GSFC, Observational Cosmology Laboratory, Code 665, USA; Bernard.J.Rauscher@nasa.gov
[2] CRESST/UMBC/GSFC,Greenbelt, MD, USA
[3] CRESST/UMd/GSFC, Greenbelt, MD, USA
[4] Sigma Space Corporation/GSFC, Greenbelt, MD, USA
[5] Markury Scientific, Inc., 518 Oakhampton Street, Thousand Oaks, CA, USA
[6] NASA/GSFC, Detector Systems Branch, Code 553, USA
[7] NASA/GSFC, Flight Software Systems Branch, Code 582, USA
Received 2017 April 26; accepted 2017 July 14; published 2017 September 6

Abstract

Near-infrared array detectors, like the *James Webb Space Telescope (JWST)* NIRSpec's Teledyne's H2RGs, often provide reference pixels and a reference output. These are used to remove correlated noise. Improved reference sampling and subtraction (IRS^2) is a statistical technique for using this reference information optimally in a least-squares sense. Compared with the traditional H2RG readout, IRS^2 uses a different clocking pattern to interleave many more reference pixels into the data than is otherwise possible. Compared with standard reference correction techniques, IRS^2 subtracts the reference pixels and reference output using a statistically optimized set of frequency-dependent weights. The benefits include somewhat lower noise variance and much less obvious correlated noise. NIRSpec's IRS^2 images are cosmetically clean, with less $1/f$ banding than in traditional data from the same system. This article describes the IRS^2 clocking pattern and presents the equations needed to use IRS^2 in systems other than NIRSpec. For NIRSpec, applying these equations is already an option in the calibration pipeline. As an aid to instrument builders, we provide our prototype IRS^2 calibration software and sample *JWST* NIRSpec data. The same techniques are applicable to other detector systems, including those based on Teledyne's H4RG arrays. The H4RG's interleaved reference pixel readout mode is effectively one IRS^2 pattern.

Key words: instrumentation: detectors – methods: statistical

Online material: color figures

1. Introduction

The Near-infrared Spectrograph (NIRSpec; Birkmann et al. 2016) is the *James Webb Space Telescope*'s (*JWST*) primary 0.6–5 μm spectrograph. NIRSpec's main mode is multi-object spectroscopy with spectral resolution $R = \lambda/\Delta\lambda = 100$, 1000, and 2700. Due to the low background provided by the observatory and the low dark current rates of the detectors, NIRSpec will be detector noise limited for most faint object observations. In this regime, the exposure time needed to achieve a given signal-to-noise ratio (S/N) for a given pixel scales directly with the read noise. Minimizing read noise is therefore key to maximizing NIRSpec's performance.

Recovering the spectrum of faint objects usually involves operations on more than one pixel. We therefore identified correlated noise ($1/f$ banding, etc.) as a second noise feature that limits NIRSpec's sensitivity. Correlated noise is particularly important for NIRSpec's multi-object spectrograph (MOS) mode because having less correlated noise enables more efficient spectral extraction. In particular, it reduces the need for local sky samples to mitigate correlated noise.

If useful local sky is always available (e.g., in a sparse deep field), then one can can achieve sensitivity within a few percent of IRS^2 in NIRSpec's traditional mode. However, this places requirements on the astronomical scene. Achieving lower correlated noise has the potential to increase NIRSpec's MOS multiplex advantage in more complex scenes by making non-local sky samples more competitive.

In this context, we developed Improved Reference Sampling and Subtraction (IRS^2) to reduce the read noise of NIRSpec's Teledyne SIDECAR™ ASIC (hereafter "SIDECAR") and H2RG-based detector system to below what is possible in *JWST*'s traditional "MULTIACCUM" readout (Rauscher et al. 2007). IRS^2 works by making more efficient use of the H2RG's

Original content from this work may be used under the terms of the Creative Commons Attribution 3.0 licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

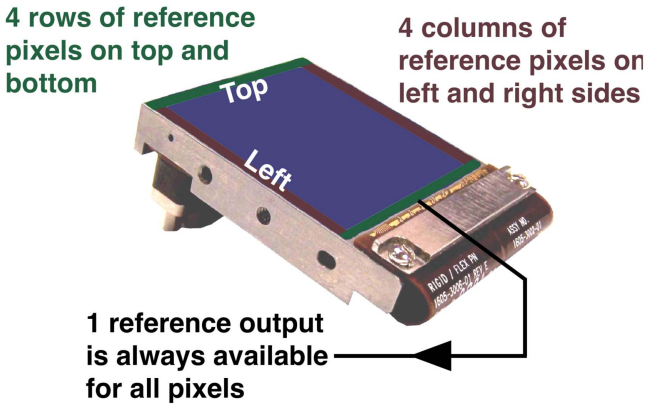


Figure 1. Teledyne's H2RG detector array provides reference pixels and a reference output. Although they do not respond to light, these references are designed to electronically mimic regular pixels. As such, they can be used to remove correlated noise. There are four rows of reference pixels along the top and bottom edges, and four columns or reference pixels on the left and right sides. The reference output is physically located at one edge of the array. It has its own video output and is available at all times for all pixels. The H2RG produces 2048 × 2048 pixel full-frame images. However, only pixels [4:2044, 4:2044] (using zero-offset Python notation) respond to light.

(A color version of this figure is available in the online journal.)

reference pixels and reference output than is possible on conventional readout schemes such as MULTIACCUM. Although we developed IRS^2 for NIRSpec's system, it is applicable to other detector types and controllers.[8]

For example, Teledyne's new H4RGs build in one IRS^2 readout pattern. In Teledyne documentation, this appears as "interleaved reference pixel readout." Taking the best advantage of the new mode requires the equations presented in Section 5. Even if a SIDECAR is not used for HxRG[9] control, IRS^2 can still be implemented. If using an H2RG, one would program the IR array controller to use the H2RG's internal shift registers as is described in Section 4. We believe that a similar approach could be taken with other HxRGs and, depending on the details, perhaps IR arrays from other vendors.

NIRSpec uses two 2048 × 2048 pixel Teledyne H2RG detectors. The HgCdTe detector material is sensitive over the 0.6–5 μm bandpass. The H2RG (Figure 1) builds in two kinds of reference pixels. Of the 2048 × 2048 pixel image area, the outer four rows and columns of pixels on all sides are reference pixels. Although they do not respond to light, the reference pixels are designed to electronically mimic a regular light-sensitive pixel. In particular, they include a dummy capacitor that simulates a regular pixel's capacitance and that is connected to the important detector substrate (DSUB) bias voltage.

Each frame of H2RG data therefore consists of the 2040 × 2040 photosensitive regular pixels surrounded by a four-pixel wide border of reference pixels on all sides. In addition to these reference pixels, there is a dedicated reference output that also mimics a regular pixel. Because the reference output is available at all times, it enables subtracting high-frequency common mode noise that is missed by the reference pixels.

Teledyne H2RGs and SIDECAR ASICs are widely used at observatories around the world. Most often, one SIDECAR is paired with each H2RG and used to read it out differentially. This is accomplished by routing the H2RG's video outputs to the SIDECAR's video inputs. The H2RG's one-reference output is typically also carried back to the SIDECAR, where it is routed in SIDECAR firmware to provide the negative input to the SIDECAR's differential video preamplifiers. The reference output is used in parallel for all n video channels, where in the H2RG the number of video channels is software selectable within $n \in \{1, 4, 32\}$. When operated this way, the reference output is subtracted with unity gain and completely open bandwidth from all n video channels within the first preamplifiers.

In contrast to the reference output, which is typically subtracted on the fly as described above (if it is used at all), the reference pixels are usually subtracted as part of the post processing to calibrate scientific data. Although each user has a preferred way of handling the reference pixels, there are some commonalities. Most often, each video output is treated separately from the others, and some combination of the reference pixels in rows is robustly averaged (to reject statistical outliers) and subtracted. Although only two of the video outputs have reference columns, these can also be used. When the reference columns are used, they are typically averaged and smoothed prior to applying a row-by-row reference correction to the image. Many groups arrive at the best recipe by iteratively working toward an algorithm that works well for their system.

IRS^2 is different. The *JWST* NIRSpec detector subsystem was designed to be highly linear. It was likewise designed to be extremely stable and repeatable, with the statistical properties of the noise being independent of time. These properties have been verified by testing. Building on this foundation, we realized that the NIRSpec detector subsystem could be well modeled as a covariance stationary linear system. Because the reference pixels and reference outputs are designed to mimic a regular pixel except insofar as its response to light, we assumed that in the absence of light, the dark regular pixels could be represented by a linear combination of the reference output and the reference pixels. Because the system is covariance stationary to a high degree of approximation, IRS^2 was developed in Fourier space. For a covariance stationary system,

[8] Better detectors and better controllers are always beneficial. We anticipate that one could do even better than was achieved in NIRSpec by using lower noise controllers and detectors together with IRS^2 to reject correlated noise irrespective of its origin in the detector or controller.

[9] We define HxRG as a generic identifier for Teledyne's H4RG, H2RG, and H1RG near-IR array detectors. Apart from pixel count, the architecture is broadly similar within the HxRG family.

Table 1
Comparing Traditional and IRS^2 Readout in *JWST* NIRSpec

Item	Traditional	IRS^2
Reference pixels in rows	Four rows of reference pixels on top and bottom of every frame.	Four rows of reference pixels on top and bottom of every frame. In addition, r reference pixels from one reference row interleaved every n normal pixels throughout the frame. For NIRSpec, the standard values are ($n = 16$, $r = 4$). The interleaved reference pixels sample both even and odd-numbered columns to allow removal of alternating column pattern noise.
Reference pixels in columns	Four columns of reference pixels in left and right outputs only.	Same.
Reference output	The reference output is subtracted in real time and with unity gain from each video output. It provides the reference for the SIDECAR ASIC's differential video input amplifiers.	The reference output is digitized in parallel with the four other video outputs and saved for later IRS^2 processing. During readout, the SIDECAR ASIC's single-ended video inputs are used. The IRS^2 processing applies an optimal set of frequency-dependent weights before the reference output is subtracted in post processing.

noise that appears correlated in the time domain (pixel space) is completely uncorrelated in Fourier space.

For gauss random read noise, IRS^2 rigorously provides the best possible reference correction if least squares is accepted as the figure of merit. This is important because it means that one knows that further reference subtraction trade studies will not provide any further benefit unless some other figure of merit is adopted. For astronomical detector systems that (1) are linear, (2) have read noise that is statistically independent of time, and (3) have Gaussian read noise after reference correction, IRS^2 provides the best possible reference correction when least squares is used as the figure of merit.

Even using these techniques, early trade studies revealed that the flight system had significant correlated noise extending from DC up to $\sim$3 kHz on the low-frequency side and at the 50 kHz Nyquist frequency. For this reason, we introduced a new clocking pattern to acquire many more reference samples than is possible in traditional H2RG readout. The IRS^2 clocking pattern, which is described in Section 4, interleaves many more reference pixels than the conventional pattern. It uses these and the reference output to remove the $\lesssim$3 kHz and 50 kHz noise (Figure 8).

To summarize, the essential elements of IRS^2 compared to traditional H2RG readout are as follows.

1. IRS^2 uses a different clocking pattern to interleave many more reference pixels into the data than is otherwise possible.
2. IRS^2 subtracts the reference pixels and reference output using an optimal set of frequency-dependent weights that represent a least-squares fit of these references to a training data set.

Table 1 briefly summarizes some of the key differences between traditional H2RG readout and IRS^2.

This article is intended to provide a concise yet reasonably complete description of IRS^2. It is the culmination of six years of work by a large team. As such, it reflects a mature understanding of what IRS^2 does and why. Readers wishing to see more of the intermediate steps and background investigations may read some of our earlier papers. These include Moseley et al. (2010), which provides the first published description of the IRS^2 concept and Rauscher et al. (2011), which describes early proof of concept testing using an engineering grade H2RG and SIDECAR ASIC. In Rauscher et al. (2012b), we presented the full set of IRS^2 equations in their current form for the first time. Finally, in Rauscher et al. (2013) we performed principal components analysis of a flight like NIRSpec system and laid the foundations for extending IRS^2 to remove the small amount of non-stationary noise that remains even after IRS^2 processing.

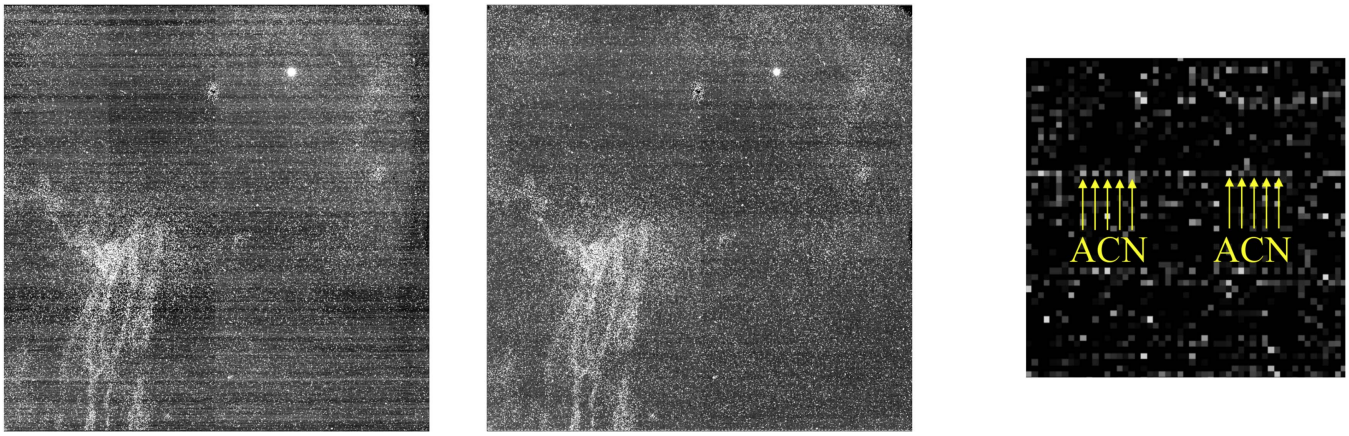
For readers who may be deciding whether to try IRS^2, Section 2 explains the benefits from a *JWST* NIRSpec perspective. Our aim is to provide enough information (between this narrative and the freely downloadable software and sample data) for readers to make an informed decision on whether to try IRS^2.

In Section 3, we explain more of the underlying physical rationale and mathematical concepts of IRS^2. This section explains why IRS^2 is expressed most naturally in Fourier space and why we believe that the same concepts are applicable to other detector systems.

In Section 4, we provide an overview of the IRS^2 implementation within the *JWST* SIDECAR assembly code. This section also explains the NIRSpec data format, which is necessary to understanding the sample data.

The *JWST* SIDECAR software source code is unfortunately controlled under the International Traffic in Arms Regulations (ITAR) and subject to other restrictions.[10] However, we are in the process of making the source code available to any United

[10] ITAR is a set of United States government regulations that pertain to specified defense-related technologies, including *JWST*'s detector systems. Under ITAR, we cannot legally publish information that would facilitate duplication of the controlled technologies.



a) Traditional b) IRS^2 c) Traditional showing ACN

Figure 2. Comparison of a (a) traditional and (b) IRS^2 readout for a pre-2010 engineering grade NIRSpec H2RG. The clouds of white warm pixels are the result of a detector degradation issue that was completely solved several years ago (Rauscher et al. 2012a). Both images are on the same [$-10\,e^-$, $+20\,e^-$] grayscale. The correlated noise (faint horizontal banding) is clearly lower in the IRS^2 image. Panel (c) shows a subarray extracted from a traditional image showing ACN. IRS^2's interleaved reference pixels mitigate ACN by including even and odd reference pixel samples.
(A color version of this figure is available in the online journal.)

States Government Agency through the NASA Technology Transfer Program. Depending on the specific circumstances, release to other U.S. persons or organizations may be possible. To provide more insight into the clocking pattern, we have written an executable Jupyter notebook in Python that is freely available for download on the *JWST* web site. Please contact the lead author for more information.

Building on Section 4, Section 5 presents the equations needed to reference-correct IRS^2 data. Of these, the most important are Equations (13), (14), and (15). These are the equations used for determining the frequency-dependent weights and applying them respectively. The other equations in this section are preliminaries to these.

Finally, we close with a summary. Readers who wish to understand IRS^2 in detail are strongly encouraged to download the IDL source code and sample data. This hands-on information provides better insight into the details than any narrative can hope to achieve.

2. Benefits and Downsides of IRS^2

For *JWST* NIRSpec, the most important benefit of IRS^2 is suppressing correlated noise. The noise appears primarily as horizontal banding oriented perpendicular to the spectral dispersion direction. This corresponds to the H2RG's fast scan direction. The faint banding is caused by $1/f$ drifts in the SIDECAR ASIC readout electronics. Another important correlated noise source is alternating column noise (ACN). The ACN is a Teledyne proprietary artifact of how the even and odd column buses are implemented in HxRG readout integrated circuits (ROIC). Figure 2 shows some examples of how correlated noise appears in traditionally sampled and IRS^2 sampled NIRSpec data.

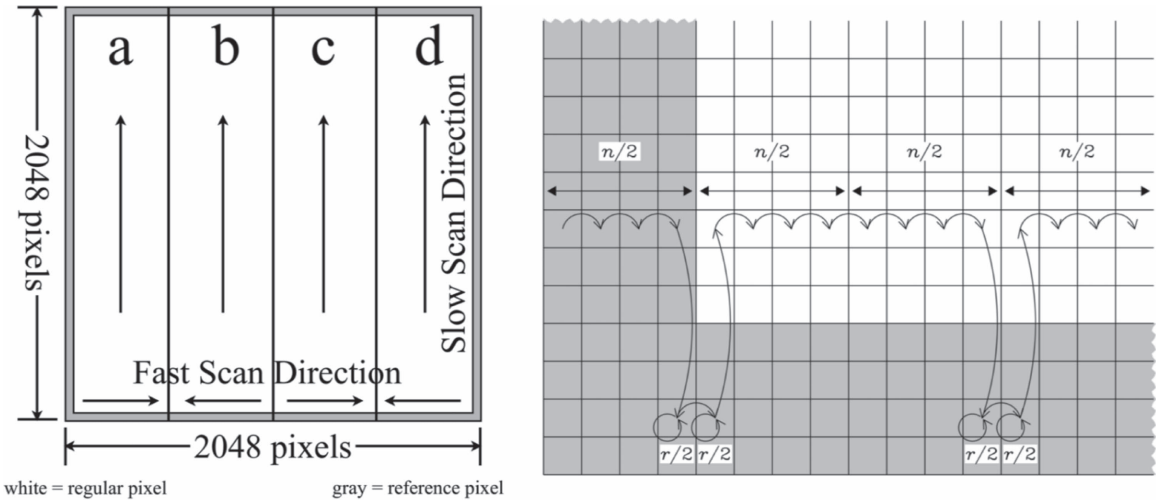
Readers interested in understanding more about H2RG read noise in general should see Rauscher (2015), which also includes a free downloadable Python language noise generator that can be used to produce realistic read noise for traditionally sampled H2RG systems.

As highlighted in the introduction, for NIRSpec the lower noise variance and reduced correlated noise translate directly into better scientific performance. For individual objects, IRS^2 provides higher S/N per unit observing time. The lower correlated noise that IRS^2 provides enables novel high multilplex observing strategies using non-local sky samples.

One might ask about the possible downsides of IRS^2. For NIRSpec, implementing IRS^2 required somewhat more development time (with associated cost) on account of the added complexity. Regarding technical downsides, IRS^2 sends more data to the ground. In the NIRSpec implementation, previously existing *JWST* data volume constraints (in other parts of the system) limited us to 65 up-the-ramp frames compared with 88 up-the-ramp frames in traditional mode. However, the reduced number of up-the-ramp frames is not fundamental to IRS^2. This could be fixed in a new system that did not have the same constraints.

3. Why IRS^2 Works

Astronomical detector systems are designed to be very stable, both in their response to light and read noise. In other words, they are designed so that to a very high degree of approximation, the read noise's covariance matrix is independent of time. The read noise is therefore very nearly covariance



a) Traditional *JWST* full-frame readout b) IRS^2 readout, lower-left corner only

Figure 3. (a) In traditional H2RG clocking, pixels are read out and digitized in the same order as they appear physically on the H2RG ROIC. For traditional *JWST* clocking, the H2RG is read out using four video outputs. These appear as thick vertical bands here. The fast scanners alternate directions and readout proceeds one row a time from bottom to top. The photosensitive area is bordered on all sides by a four-pixel wide frame of reference pixels. (b) zooms in on the lower left-hand corner of the H2RG to show how the IRS^2 clocking pattern differs to interleave reference and normal pixels. For every n normal pixel, we step out to a reference row to digitize an even-numbered reference pixel $r/2$ times and then an odd-numbered reference pixel $r/2$ times. We sample both even and odd reference columns to subtract alternating column pattern noise at the Nyquist frequency (Moseley et al. 2010). In IRS^2, the reference output is digitized and saved in parallel with the four video outputs. The H2RG's vertical window mode scanner is used to implement stepping out to the reference pixels, which is effectively no different from reading a one-pixel subarray that happens to be in the same column as the science pixel. This figure omits some timing overheads for clarity. See Appendix B for more information on the timing, including an explicit discussion of the NIRSpec timing overheads.

stationary. It is a classical result that the eigenspace of a stationary covariance matrix is Fourier space. The Fourier basis vectors provide an orthogonal representation of the read noise.

This is important because when stationary noise is viewed in Fourier space, the different frequencies are uncorrelated. Even though the data may contain strong pixel-to-pixel correlations when viewed in the time domain (as images), there are no correlations between frequencies when the same data are viewed in Fourier space.

With this realization, we understood that we could treat each frequency independently to model the regular pixels as a linear combination of all available reference information. In the current implementation, the model includes the reference pixels (including additional reference samples) and the reference output. If other reference timeseries were to become available, we could easily modify the IRS^2 equations to include them.

4. *JWST* IRS^2 SIDECAR Assembly Code Implementation

4.1. Simple Prototype Implementation

Compared with the traditional full-frame readout pattern, the IRS^2 clocking pattern contains two distinct differences. First is the interleaving of H2RG reference pixels (from one specified reference row) within the science data stream. This is shown in Figure 3. Second is the separate digitization of the H2RG reference output (i.e., single-ended readout) to facilitate differential readout in ground post processing instead of in the SIDECAR pre-amplifier block (Figures 4–5).

"Stepping out" from the regular pixels to the reference row and "stepping in" again is accomplished using the H2RG's movable guide window. In effect, we program a movable one-pixel guide window in a pre-selected reference row that shadows the regular pixels. The full-frame horizontal scanner is always used, while the full-frame vertical scanner is used for regular pixels and the guide window vertical scanner is used for the reference row.

We began developing IRS^2 at a time when *JWST* flight software development was already well underway. To implement this concept within the existing *JWST* flight system, multiple software components needed to be enhanced including the SIDECAR assembly code and the *JWST* Integrated Science Instrument Module's (ISIM) Command and Data Handling Software (ICDH). Some of these enhancements would be applicable to any HxRG/SIDECAR system, and others are only necessary to work around existing downstream hardware limitations on *JWST*.

The *JWST* prototype and flight implementation allows for user configurable parameters to specify the details of the readout (n = number of normal science pixels and r = the number of interleaved reference pixels) and the row within the H2RG to use as reference pixels. Figure 3(b) shows how n and r appear in the

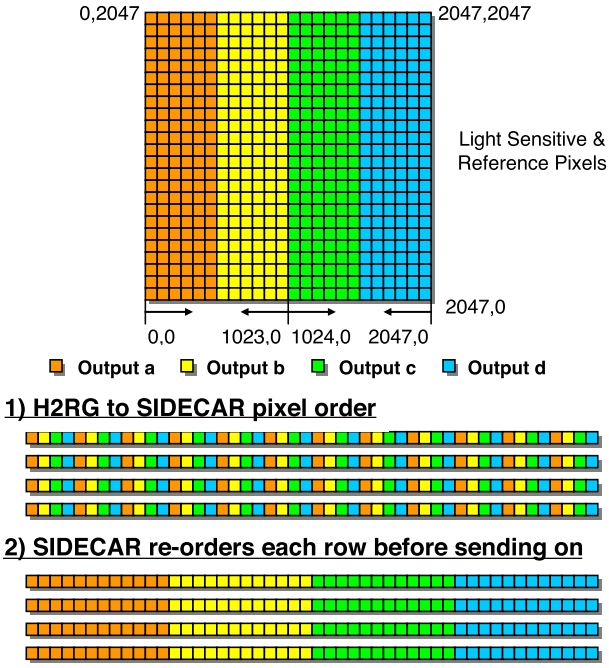


Figure 4. Traditional full-frame H2RG readout in which the reference output is subtracted in real time in the SIDECAR's video preamplifiers, hence the reference output does not appear here. (1) Outputs a, b, c, and d are sampled and digitized simultaneously. (2) The outputs are de-interlaced in the SIDECAR prior to being sent on for further processing. This figure contains some of the same information as Figure 3(a). We include it because it highlights how the data are transmitted in the flight data systems. This will serve as a useful point of comparison to Figure 5, which does the same for IRS^2 readout.

(A color version of this figure is available in the online journal.)

timing pattern. The reference row can be selected from the bottom or top four rows, all of which are reference rows. Through experimentation with both the prototype (engineering grade) and NIRSpec flight hardware, ($n = 16$, $r = 4$) was determined to be near optimal for NIRSpec. In practice, we found that it did not matter greatly which reference row was selected.

Implementation of the interleaved pixel readout on the *JWST* hardware utilizes the H2RG's vertical window mode scanner, which was originally intended for support of guide mode. Before the exposure begins, the SIDECAR assembly code positions the vertical window mode scanner at the user specified row within the H2RG as part of preparation of the acquisition. During readout, the transition from the science pixels to the reference pixels is done by selecting the appropriate vertical scanner. The single full-frame horizontal scanner is utilized for both science and reference pixel readouts. The time needed to transition between vertical scanners injects a single pixel time gap between the sampling of the science and reference pixels (Figure 6).

Finally, to implement the single-ended digitization of the H2RG reference output, the signal is routed via the SIDECAR's global routing bus into a dedicated video channel in parallel with the other four video channels. In contrast with traditional readout, the five video channels are configured for single-ended operation.

This implementation contains all readout changes needed to implement IRS^2 on standard SIDECAR + H2RG system as was done with the *JWST* prototype code, which operates with a JADE2, and free of constraints from the other flight elements of *JWST*. For most readers, the prototype code will be the most straightforward to adapt to their own systems.

4.2. Additional Challenges in Flight Implementation

As discussed, the *JWST* flight implementation had further complexities due to the challenges of implementing IRS^2 in the already-built flight systems. Creative use of the values within the SIDECAR science packet headers were required to have the data flow properly to the downstream electronics. Each row of the H2RG was split in half for compatibility with the NIRSpec focal plane electronics, and each frame was also split in half to accommodate the memory limitations within the data system memory. Finally, the pixel ordering within the flight code needed to be updated significantly for IRS^2, again to work properly with the downstream electronics design. If a reader chooses to review the *JWST* flight code, then these additional complexities should be accounted for, but are unlikely to manifest in other systems.

5. Reference Correcting IRS^2 Data

After acquiring the data, IRS^2 reference correction should be the first step in the calibration process. Once the raw data cubes have been reference corrected, subsequent calibration steps including linearity correction and up-the-ramp slope fitting can be done using the standard tools that are available at most observatories.

For NIRSpec, we found that it is not possible to improve upon IRS^2 reference correction with any further reference pixel correction. As expected, subsequent use of the reference rows and columns degrades the noise by at least a few percent. If attempted, the degradation typically appears as a small increase in the correlated noise.

The H2RG provides three types of output: n_p normal pixels, r_p reference output, and ρ_p reference pixels. In IRS^2, these are represented by vectors where p is a time domain index that runs over all time steps in the exposure (and incidentally over all pixels because they are time ordered). In Moseley et al. (2010), we showed that the read noise's eigenspace is close to Fourier space. We therefore work in Fourier space because the basis vectors (sines and cosines) are linearly independent. Fourier transforming each output, we arrive at n_ν, r_ν, and ρ_ν, where ν is is an index that runs over frequency.

IRS^2 is a linear model. The reference corrected normal pixels are represented by

$$n'_\nu = n_\nu - \alpha_\nu r_\nu - \beta_\nu \rho_\nu, \quad (1)$$

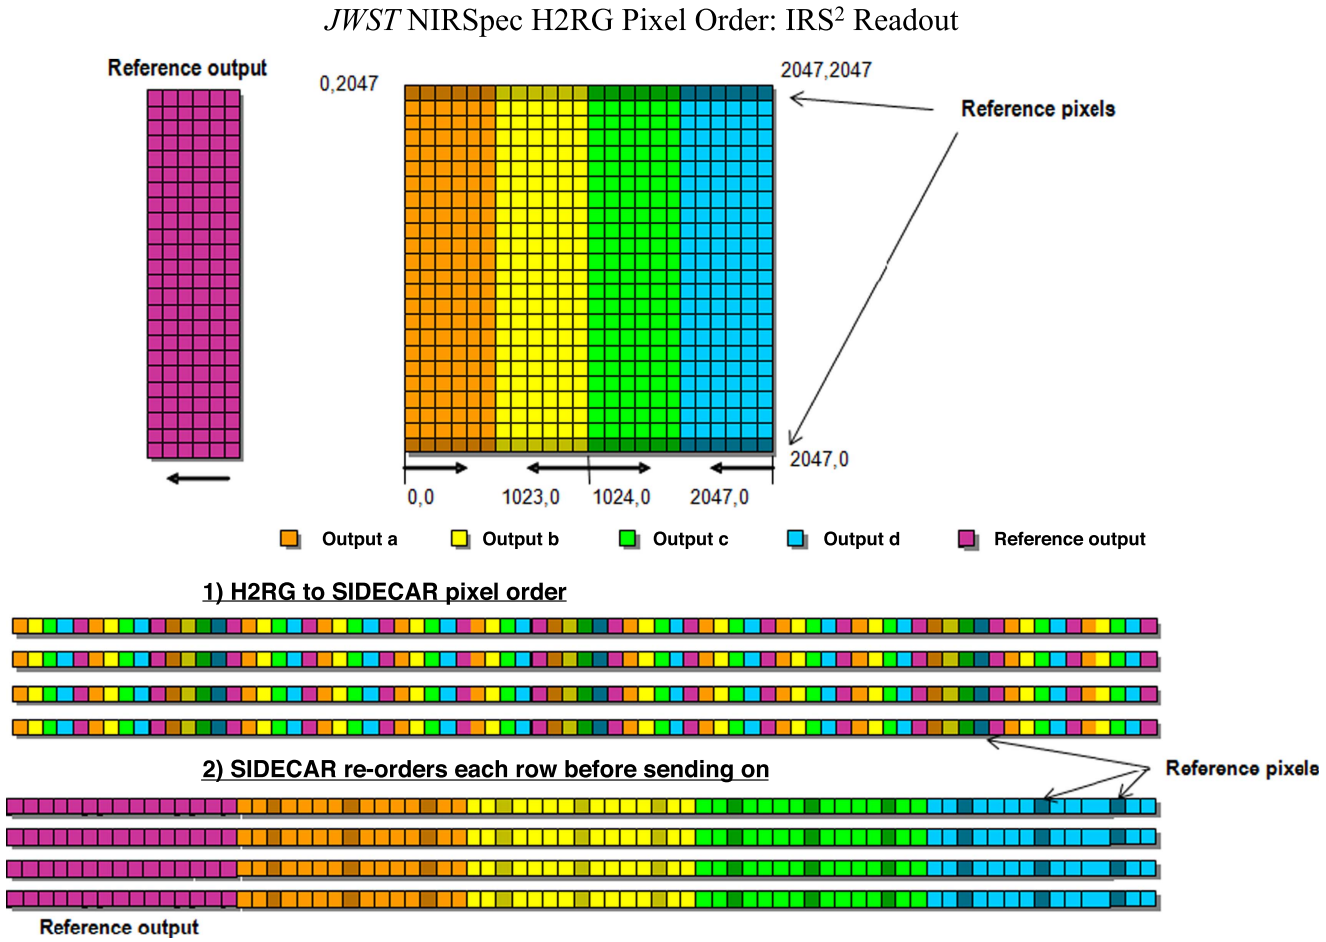


Figure 5. This figure should be viewed in contrast with Figure 4. In IRS^2, shown here, the reference output is sampled and digitized in parallel with the regular outputs a, b, c, and d. These are de-interlaced in the SIDECAR as before prior to being sent for further processing. The dark shading indicates reference pixels. The H2RG has four rows of reference pixels on the top and bottom edges. Here, we highlight all reference rows with dark shading. In practice, only one of the eight reference rows is selected and used to provide all interleaved reference pixels.
(A color version of this figure is available in the online journal.)

where α_ν and β_ν are vectors of frequency-dependent weights that function analogously to Wiener filters.

If we have a sufficiently large number of training dark frames (typically $\geqslant 10^3$ for NIRSpec), and i is an index that runs over this set, then we can use the method of least squares to solve for α_ν and β_ν. Let

$$X_\nu^2 = \sum_i (n_\nu^i - \alpha_\nu r_\nu^i - \beta_\nu \rho_\nu^i)(n_\nu^{i*} - \alpha_\nu^* r_\nu^{i*} - \beta_\nu^* \rho_\nu^{i*}) \quad (2)$$

be the least-squares figure of merit, where * denotes the complex conjugate. X_ν^2 is minimized when

$$\alpha_\nu = \frac{\left(\sum_i n_\nu^i r_\nu^{i*}\right)\left(\sum_i \rho_\nu^i \rho_\nu^{i*}\right) - \left(\sum_i n_\nu^i \rho_\nu^{i*}\right)\left(\sum_i \rho_\nu^i r_\nu^{i*}\right)}{\left(\sum_i r_\nu^i r_\nu^{i*}\right)\left(\sum_i \rho_\nu^i \rho_\nu^{i*}\right) - \left(\sum_i \rho_\nu^i r_\nu^{i*}\right)\left(\sum_i r_\nu^i \rho_\nu^{i*}\right)} \quad (3)$$

and

$$\beta_\nu = \frac{\left(\sum_i n_\nu^i \rho_\nu^{i*}\right)\left(\sum_i r_\nu^i r_\nu^{i*}\right) - \left(\sum_i n_\nu^i r_\nu^{i*}\right)\left(\sum_i r_\nu^i \rho_\nu^{i*}\right)}{\left(\sum_i r_\nu^i r_\nu^{i*}\right)\left(\sum_i \rho_\nu^i \rho_\nu^{i*}\right) - \left(\sum_i \rho_\nu^i r_\nu^{i*}\right)\left(\sum_i r_\nu^i \rho_\nu^{i*}\right)}. \quad (4)$$

One can factor Equations (3)–(4) into a convenient set of sums that can be augmented whenever new darks become available to improve the coefficients.[11] These are as follows:

$$N_\nu = \sum_i n_\nu^i n_\nu^{i*}, \quad (5)$$

$$R_\nu = \sum_i r_\nu^i r_\nu^{i*}, \quad (6)$$

[11] If a substantial change is made to the system (e.g., changing a bias voltage or the operating temperature), we recommend recalibrating the IRS^2 coefficients using new training data.

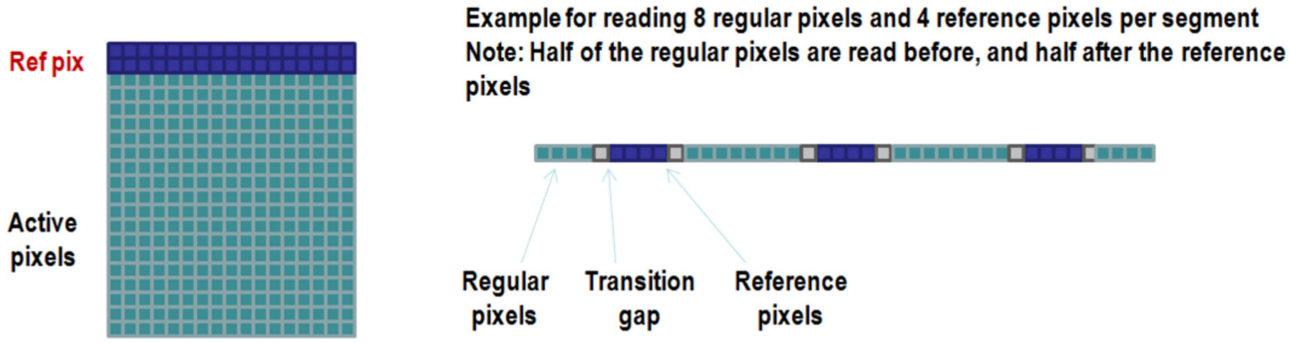


Figure 6. There is a one-pixel time "transition gap" associated with stepping out to the reference pixels or stepping back into the science pixels. (A color version of this figure is available in the online journal.)

$$P_\nu = \sum_i \rho_\nu^i \rho_\nu^{i*}, \tag{7}$$

$$X_\nu = \sum_i n_\nu^i \rho_\nu^{i*}, \tag{8}$$

$$Y_\nu = \sum_i n_\nu^i r_\nu^{i*}, \tag{9}$$

$$Z_\nu = \sum_i \rho_\nu^i r_\nu^{i*}. \tag{10}$$

The vectors N_ν, R_ν, and P_ν are real, while X_ν, Y_ν, and Z_ν are complex.

With these definitions, we can rewrite Equations (3)–(4) as

$$\alpha_\nu = \frac{YP - XZ}{RP - ZZ^*}, \text{ and} \tag{11}$$

$$\beta_\nu = \frac{XR - YZ^*}{RP - ZZ^*}, \tag{12}$$

where we have suppressed the ν suffix on the right-hand side to achieve a more compact notation. In Equations (11)–(12) the denominators are real, but the numerators are in general complex. In general, α_ν and β_ν are therefore complex.

At intermediate frequencies, the interleaved reference pixels vector ρ_ν, is not sampled. We therefore tailor β_ν with an apodized filter, f_ν. Figure 7 shows the filter that we used while developing the standard NIRSpec IRS^2 ($n = 16$, $r = 4$) clocking pattern. Appendix B provides more information on how the filter was designed and specifically on how the roll off frequency was chosen.

Including the filter, Equations (11) and (12) can be rewritten as

$$\alpha_\nu = \frac{Y - ff^*XZ/P}{R - ff^*ZZ^*/P}, \text{ and} \tag{13}$$

$$\beta_\nu = f^* \frac{XR/P - YZ^*/P}{R - ff^*ZZ^*/P}. \tag{14}$$

Once the coefficients have been determined using Equations (13) and (14), Equation (1) simplifies to

$$n_p' = n_p - \mathcal{F}^{-1}[\alpha\mathcal{F}(r_p) + \beta\mathcal{F}(\rho_p)], \tag{15}$$

were $\mathcal{F}$ is the familiar discrete Fourier Transform. There is no need to Fourier transform the normal pixels in the pipeline. Moreover, one could implement these operations as convolutions in the time domain if desired.

This concludes the core set of equations that are needed to implement IRS^2. The frequency-dependent weights are inferred from a set of training darks using Equations (13) and (14). Once these weights are known, Equation (15) is used to apply them to the data.

5.1. Understanding α_ν and β_ν

Recall that β_ν is a set of frequency-dependent weights for the interleaved reference pixels and α_ν is a similar vector for the reference output. Figure 8 shows the measured values for a prototype implementation of IRS^2 (i.e., not the flight system).

Figure 8(a) plots all frequencies and shows that the interleaved reference pixels primarily correct very low frequencies, including $1/f$, and ACN at 50 kHz. The low-frequency wing that is visible near 50 kHz is caused by column switching acting as a carrier for much lower frequency $1/f$. The amplitude of β_ν is zero at intermediate frequencies, as expected given that these frequencies have been filtered out.

The reference output, α_ν, has amplitude at all frequencies, but it too is strongest at low frequency. Figure 8(b) highlights the low frequencies. This figure shows one of the more important early findings in IRS^2 development. The reference output does not have unity gain at low frequency. The behavior shown here is fairly typical for NIRSpec H2RGs. The reference output should be subtracted with less than unity gain. In traditional *JWST* readout, it is subtracted with unity gain and this explains why it is not more effective at subtracting out $1/f$ noise in traditional readout.

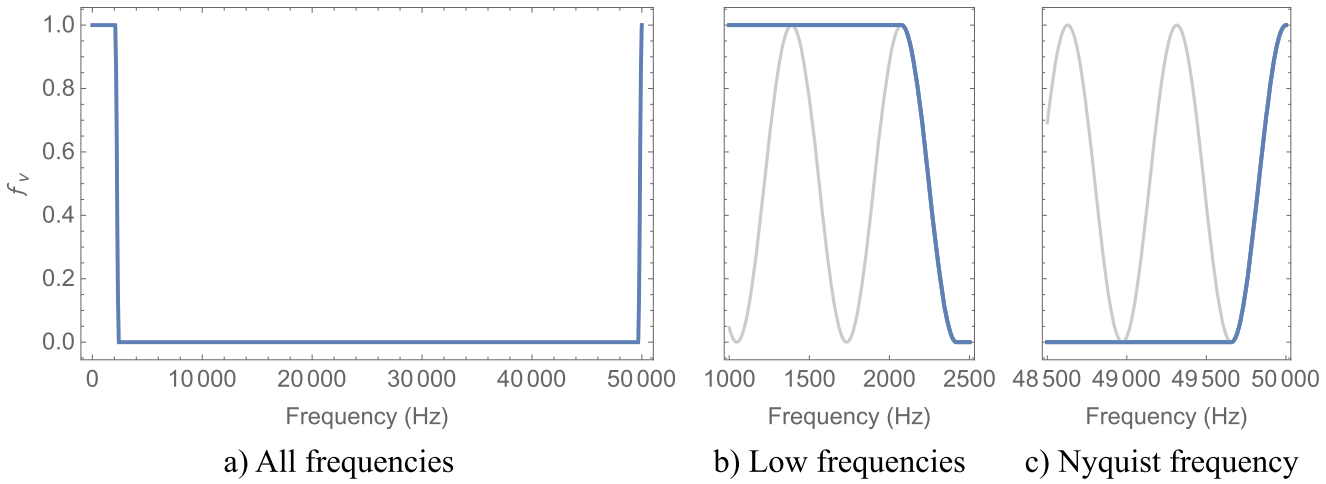


Figure 7. Interleaved reference pixels, ρ_ν, sample only low-frequency noise and alternating column noise at the Nyquist frequency. To avoid injecting mid-frequency noise, we apply an apodized filter, f_ν. Panel (a) shows the filter that is used with the standard NIRSpec ($n = 16$, $r = 4$) IRS^2 clocking pattern. In panel (b), we use one quadrant of a sine function to roll the response from unity to zero with the half-power point at $f_{1/2} = 2,247.19$ Hz. The gray curve is a graphical element to highlight the sinusoidal shape of the roll off. Panel (c) is mirrored at the Nyquist frequency to enable IRS^2 removal of alternating column pattern noise. The width of the roll off and roll on is somewhat arbitrary at present. We empirically determined that 5000 frequency steps (=342.894 Hz) works well for NIRSpec. Appendix B explains more fully how we determined the half-power point for this particular IRS^2 pattern.
(A color version of this figure is available in the online journal.)

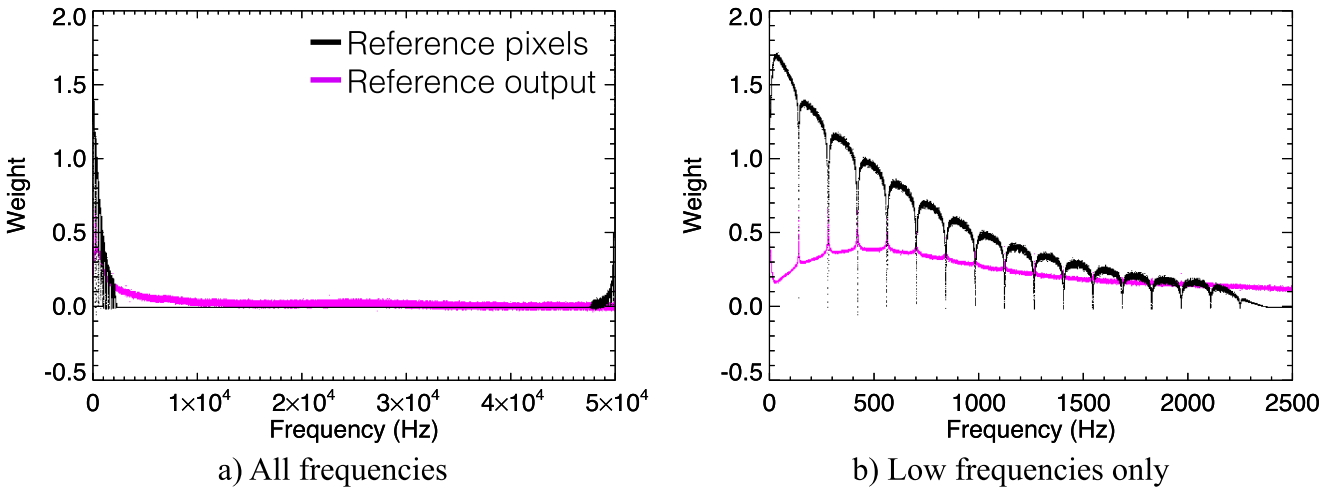


Figure 8. In IRS^2, the reference output and interleaved reference pixels are filtered with a set of frequency-dependent weights. Panel (a) plots all frequencies for a non-flight prototype of the NIRSpec detector subsystem. It shows that the reference output, α_ν, is important at all frequencies and largest at low frequencies. The interleaved reference pixels, β_ν, have power only at very low frequencies and near the Nyquist frequency. There is a wing on the low-frequency side of 50 kHz. This is caused by column switching acting as a carrier for much lower frequency $1/f$. Panel (b) highlights the low frequencies. Interestingly, the reference pixels have gain <1 at all frequencies for this system. This partially explains why traditional *JWST* readout, which subtracts the reference output with unity gain, is not optimal.
(A color version of this figure is available in the online journal.)

5.2. Optimal Use of the Reference Output when Interleaved Reference Pixels are not Available

For *JWST*, we were in the fortunate position of being able to write our own SIDECAR ASIC software. This allowed us to implement the interleaved reference pixels that are a hallmark of IRS^2. For groups who are not in a position to do this, the IRS^2 study nevertheless provided insight for how to best use the reference output.

If only the reference output is used to read an H2RG differentially, then α_ν needs to be recalculated. Under these

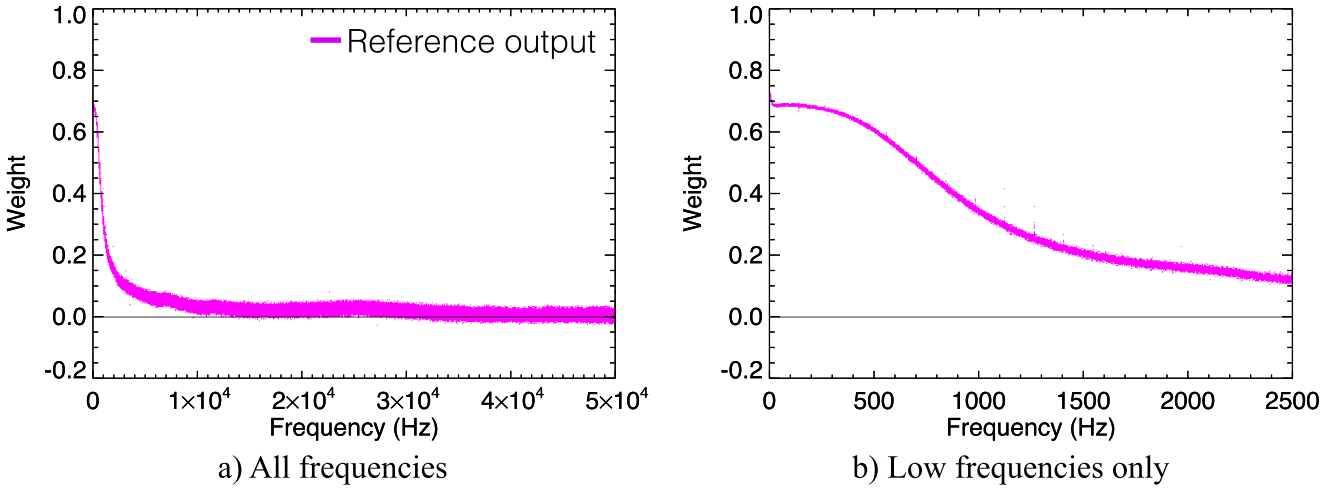


Figure 9. If only the reference output is used to make the correction, then α_ν needs to be recomputed. Here, we show the result of recalculating α_ν using the same training data as were used to produce Figure 8. The weight is less that unity at low frequencies. As expected, there is no weight at 50 kHz because the reference output does not experience column switching. Panel (b) suggests a potentially simple way of improving the performance of existing systems by filtering the reference output and operating differentially. This is described more fully in the text.
(A color version of this figure is available in the online journal.)

conditions, α_ν takes all of the weight and is much closer to 1.0 at low frequencies (typically within 20%), although α_ν still falls off at higher frequencies. Moreover, it contains no trace of any correction for the ACN because the reference output does not see the even and odd columns. Figure 9 shows an example of recalculating α_ν only using the same input data as were used in Figure 8.

The considerable change at low frequency is because r_ν and ρ_ν are highly correlated at low frequencies, thus either can be used to remove $1/f$ noise. When only the reference output is available, it takes all of the weight, therefore α_ν approaches unity. The results shown in Figure 8 indicate that for this particular array, the interleaved reference pixels happened to be more effective at removing the lowest frequency noise.

We note in passing that Figure 9(b) suggests that some HxRG systems could potentially be improved incrementally by applying a simple passive low-pass filter to the reference output and operating the detector differentially. Although we have not yet experimentally tested this concept, we plan to do so in the near future. We do not expect filtering the reference output alone to be as powerful as IRS^2 (for example, it does nothing about ACN), but it may provide a simple way to boost the performance of existing systems without writing new controller software.

5.3. Other Implementation Details

Because IRS^2 is a linear model, a full implementation requires interpolating over all gaps and other non-linear events in the timeseries before applying these equations (e.g., overheads at the ends of rows and frames as well as "bad pixels" and cosmic ray hits). In early prototypes, we used simple linear interpolation for this. We now use the somewhat more sophisticated interpolation scheme that can be seen in the source code. The interested reader is referred to the IDL source code that is available at http://jwst.nasa.gov/publications.html for these implementation details.

For completeness, we note that if additional reference information that tracks the stationary read noise were to become available, then IRS^2 could straightforwardly be extended to incorporate it. In this case, one would revise Equation (1) to include the new reference vector and update the other equations accordingly.

6. Conclusion

IRS^2 is a technique for reducing the correlated read noise of near-IR detector systems. IRS^2 was conceived, implemented, and tested by the *JWST* Project at NASA Goddard for NIRSpec, which uses Teledyne H2RGs and SIDECAR ASICs. Compared with traditional HxRG readout, IRS^2 uses a new clocking pattern to interleave many more reference pixels into the data than is otherwise possible. As part of the IRS^2 post processing, IRS^2 subtracts both the reference pixels and reference output using a set of least squares optimized frequency-dependent weights. These weights were measured for the as-built hardware using the equations in Section 5 to least-squares fit the references to an extensive training data set.

For NIRSpec, IRS^2's primary benefit is to significantly reduce correlated noise. Compared with traditional *JWST* readout, IRS^2 images are cosmetically cleaner, with fewer instrument signatures (less $1/f$ banding, less alternating

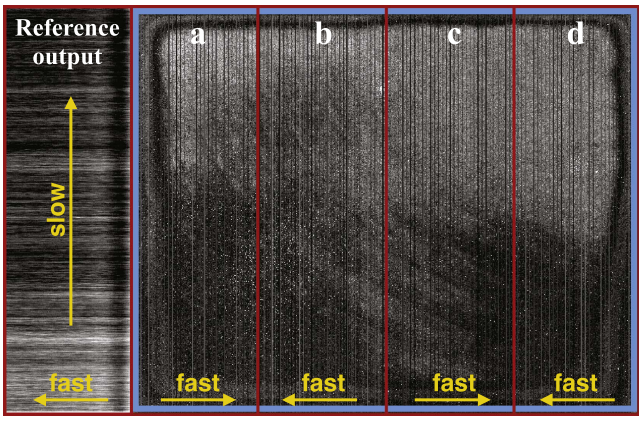


Figure 10. *JWST* NIRSpec IRS^2 data format for one frame of H2RG data. To more clearly show fine structure, we have zero offset and renormalized each output. The fast and slow scanner directions are shown in yellow. The usual four-pixel-wide border of reference pixels that surrounds each frame of H2RG data is shown in blue. Data from the reference output appears at left. The four regular outputs are labeled a, b, c, and d. Each output is wider than 512 pixels because it includes regular pixels and the extra interleaved reference pixels. As in traditional readout, each output is 2048 pixels high. If one looks carefully at any of the outputs a–d, there are faint vertical stripes where the interleaved reference pixels appear at a lower signal level than the regular pixels on the same output.

(A color version of this figure is available in the online journal.)

column noise, etc). The cosmetically cleaner images will allow the use of more distant sky samples, thereby increasing MOS multiplex advantage in crowded fields. The cosmetically cleaner images should likewise increase the efficiency of integral field unit (IFU) observations by increasing the S/N for realistic sky subtraction scenarios of extended sources. Our simulations suggest that S/N gains as great as 45% per unit observing time are potentially possible depending on the source. In any case, for a read noise limited instrument like NIRSpec, lower noise is always better. We recommend that readers who are interested in exploring noise trades download the sample data because the results depend critically on the observing scenario.

As an aid to groups who may wish to explore the benefits of IRS^2, we are making our prototype IRS^2 calibration software and sample *JWST* NIRSpec data freely available for download, which can be found on the JWST web site. Although the SIDECAR ASIC detector readout software are ITAR sensitive (and subject to other controls), they are nevertheless available to other United States Government Agencies. Depending on the circumstances, we may be able to release the SIDECAR software to other United States entities. Please contact the lead author for more information.

Looking to the future, Teledyne coordinated with us while developing the H4RG series of near-IR detector arrays. The H4RGs build in one IRS^2 readout pattern. Teledyne refers to this as "interleaved reference pixel readout." If the new mode works as hoped, this will free users from having to write the kind of IRS^2 clocking software described in Section 4. In the near future, we plan to explore the new readout mode using the techniques and equations described in Section 5.

Looking still further into the future, we are eager to explore the utility of blanked off regular pixels in NIRSpec scenes by extending IRS^2 to include them. In principle, blanked off regular pixels are an even better proxy of regular pixels than reference pixels. We plan to do this using a similar least-squares approach once appropriate NIRSpec science data start to become available in mid to late 2019.

This work was supported by NASA as part of the *James Webb Space Telescope* Project. The authors thank Drs. Pierre Ferruit and Marco Sirianni of the European Space Agency for carefully reading the entire manuscript and making several very helpful suggestions. We also thank the referee for several useful comments and helpful ideas for providing more insight into the SIDECAR clocking software within the constraints of the ITAR.

Appendix A
JWST NIRSpec IRS^2 Data Format

This appendix describes the prototype data format that was used to generate the downloadable sample data. The flight format may differ. Please consult the NIRSpec instrument documentation at the Space Telescope Science Institute (STScI) for information on the flight data format.

The traditional *JWST* clocking pattern reads the H2RG detectors using four outputs. The resulting 2048 × 2048 pixel frames of data appear in thick 2048 × 512 pixel stripes (Figure 3). Because IRS^2 returns the reference output and the four regular outputs with additional reference pixels interleaved, the resulting frame format is different.

Figure 10 shows one frame of IRS^2 sampled data. In the default ($n = 16$, $r = 4$) IRS^2 configuration, the resulting frame size is 2048 × 3200 pixels.

Appendix B
More Information on Pixel Timing and Filtering Interleaved Reference Pixels

JWST's H2RG detectors use a 100 kHz pixel clock, resulting in a 50 kHz Nyquist frequency. Early IRS^2 studies (Moseley et al. 2010; Rauscher et al. 2011) showed that the reference pixels were correlated with the regular pixels for frequencies lower than (very roughly) 3 kHz, and again near 50 kHz. At intermediate frequencies, the reference pixels do not correlate well with the regular pixels. Informed by this and based on practical considerations for implementing IRS^2 in the already-built *JWST* data systems, we selected ($n = 16$, $r = 4$) as the standard NIRSpec IRS^2 clocking pattern. Figure 11 presents the timing sequence for one row of one output on a timeline.

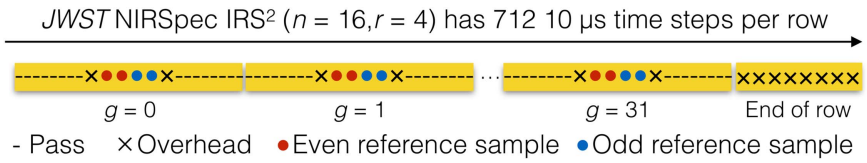


Figure 11. Timeline of the standard NIRSpec ($n = 16$, $r = 4$) IRS^2 clocking pattern for one row of one video output. There are 22 10 μs time steps within each IRS^2 group, g. When sampling regular pixels, the reference vector passes over them, leaving gaps in the reference pixel timeseries. Other gaps include a one-pixel time overhead for stepping out from the regular pixels to the reference pixels, and another one-pixel overhead when stepping back into the regular pixels. At the end of each row, there is an eight-pixel time overhead for starting the next row. The corresponding row rate is 140.449 Hz. We selected this pattern to sample (mostly) $1/f$ noise at frequencies $\lesssim$3 kHz and alternating column pattern noise at the 50 kHz Nyquist frequency. We were also required to select values of n and r that were practical to implement in the already-built *JWST* data systems.
(A color version of this figure is available in the online journal.)

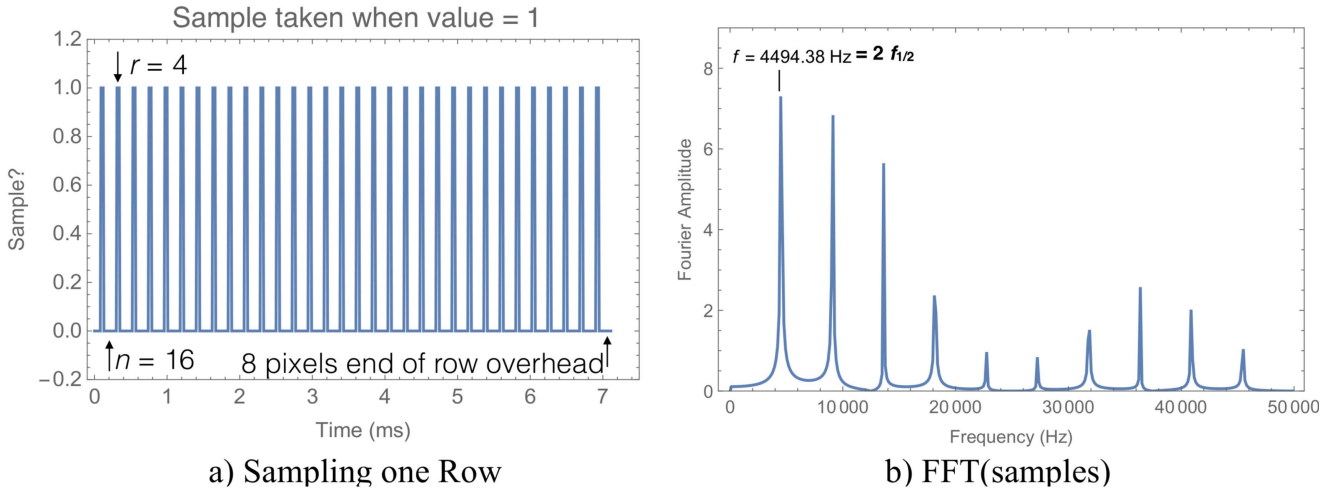


Figure 12. In NIRSpec IRS^2 ($n = 16$, $r = 4$) readout at the 100 kHz pixel rate, the sampling interval for the interleaved reference pixels is $22 \times \tau_{\rm pixel} = 220$ μs. The corresponding sampling frequency is 4,545.45 Hz. However, (a) because of the eight-pixel end of row overhead, the samples are not equally spaced in time. To compute an effective Nyquist frequency for the unequal sampling, we (b) compute the FFT of the sampling function (panel a) and measure the location of the first maximum. The half-power point of the filter shown in Figure 7 is set to one half this frequency, $f_{1/2} = 2,247.19$ Hz. In practice, this is only slightly different than the more usual "half the sampling frequency" rule, which would have worked about as well.
(A color version of this figure is available in the online journal.)

In the standard ($n = 16$, $r = 4$) IRS^2 pattern, $r = 4$ reference samples are interleaved for every $n = 16$ normal pixels. The pattern acquires two samples of a reference pixel in an even-numbered column and two samples of the next reference pixel, which is in an odd-numbered column. The even and odd reference samples are taken to remove ACN. ACN originates in a well understood, but Teledyne proprietary aspect of how the H2RG columns are biased in ROIC.

This pattern leaves gaps in the reference pixel time sequence that must be interpolated over. Simple linear interpolation will work, although we are now using the more sophisticated interpolation scheme that can be seen in the downloadable IDL code. The gaps include time for passing over the regular pixels, and also one-pixel overheads for stepping out from the regular pixels to the reference pixel and stepping in again. At the end of each row, there is an eight-pixel time overhead for starting the next row. After accounting for gaps, in each 22-pixel group, only the four reference pixel samples are saved, the first two in an even column and the second two in the next reference column.

Because the reference pixels do not correlate with the normal pixels at intermediate frequencies, IRS^2 filters out the intermediate frequencies to avoid adding noise. Figure 7 shows this filter. Here, we provide more detail on why the half-power frequency was set to $f_{1/2} = 2,247.19$ Hz for ($n = 16$, $r = 4$). This represents the (loosely speaking) Nyquist frequency of the sampling pattern including the effects of gaps.

Figure 12 explains how the half-power frequency was calculated. In panel (a), we define a reference pixels vector that includes all time steps within one row. The value was set $= 0$ when passing over a normal pixel or clocking overheads. For the four reference pixel samples in each group, the value was set $= 1$. Panel (b) shows the FFT of this vector. The maximum value occurs at $2f_{1/2} =$ Hz.

References

Birkmann, S. M., Ferruit, P., Rawle, T., et al. 2016, Proc. SPIE, 9904, 99040B
Moseley, S. H., Arendt, R. G., Fixsen, D. J., et al. 2010, Proc. SPIE, 7742, 36
Rauscher, B. J. 2015, PASP, 127, 1144
Rauscher, B. J., Fox, O., Ferruit, P., et al. 2007 PASP, 119, 768
Rauscher, B. J., Arendt, R. G., Fixen, D. J., et al. 2011, Proc. SPIE, 8155, 45
Rauscher, B. J., Hill, R. J., Greenhouse, M., et al. 2012a, AIP Adv., 2, 021901
Rauscher, B. J., Arendt, R. G., Fixsen, D. J., et al. 2012b, Proc. SPIE, 8453, 84531F
Rauscher, B. J., Arendt, R. G., Fixsen, D. J., et al. 2013, Proc. SPIE, 8860, 886005