Nancy Grace Roman Space Telescope
Technical Report Roman–STScI–000000,
SE-01

# Quantifying the Science Performance of Multi-Accumulation (MA) Tables

Timothy Brandt
January 30, 2025

## Abstract

*The Nancy Grace Roman Space Telescope has the capacity to generate huge amounts of data, but faces limitations on downlink bandwidth that preclude transmitting all of that information to Earth. Multiple detector readouts are combined on-board into resultants to reduce the data volume; this comes at some cost to the data quality. The choice of which reads to average and how is defined by multi-accumulation (MA) tables; the MA table determines how much of the information is lost in the process of on-board averaging. This technical document presents a framework and Jupyter notebook for comparing the science performance of different MA tables in order to inform their development and optimization. Any two MA tables can now be quantitatively compared as a function of a source's count rate. The Roman flight software permits a limited form of weighted averaging of the reads that combine to form resultants; this document shows that this enables further improvements in data quality and/or reductions in downlink bandwidth requirements.*

## 1 Introduction

The Nancy Grace Roman Space Telescope Wide Field Instrument (WFI) has eighteen $4096 \times 4096$ H4RG-10 detectors filling its focal plane, for a total of $3 \times 10^8$ pixels. Each detector will be read out simultaneously with 32 output channels, each of size $128 \times 4096$ pixels. At a pixel rate of ≈200 kHz this allows Roman-WFI to be fully read out nondestructively every ≈3 seconds. By saving and using every read, we can achieve the lowest read noise and effectively reject cosmic rays, which manifest as "jumps" in a pixel's charge. Unfortunately, even the 11 Tbits/day available to the Roman Space Telescope are insufficient to save every read from every WFI detector if they are being read out continuously.

Given the limitations on data transfer, Roman will combine some of its reads onboard into "resultants" and downlink these resultants; each resultant will be an average of one

or more reads. Read noise can still be mitigated and cosmic rays identified and removed, though there will be a cost in scientific data quality. The instructions for how to combine reads into resultants, called a multi-accumulation (MA) table, determine how closely the downlinked data can approach the quality achievable with all of the reads.

This technical document provides metrics for evaluating the quality of an MA table, intuition for the relative performance of different MA tables, and a description of software to compare the performance of any set of MA tables. MA tables can achieve even better performance if they are permitted to construct weighted averages of reads. This variant is described theoretically in the Appendix of Brandt (2024a). This document presents one such case as an example, together with a set of intra-resultant weights that is possible to implement within the confines of the Roman flight software.

## 2 Optimal Weights for All Reads

The first task in processing data from a detector read out nondestructively is to fit a line to the accumulated counts: to turn the raw reads into count rates at each pixel. This can be done optimally using the likelihood function assuming an ideal detector with known gain and read noise (Brandt 2024a). The result is a weighted sum of the individual detector reads. We begin by showing the optimal weights in the case of saving all of the reads, assuming that there is no cosmic ray during the ramp and that the pixel remains unsaturated at the end of the exposure. This will provide intuition for the relative performance of different MA tables, and will show the basic principles that govern optimal weighting within a resultant. For our example, we take a ramp composed of 19 reads. We fix the read noise and use a range of S/N values in the final ramp slope. The optimal weights are computed using Brandt (2024a).

In the photon noise limit (high $S/N$), all photons should be treated equally and only the first and the last reads carry any weight. In the read noise limit ($S/N \to 0$) the first and last reads still carry the largest magnitude weights, with the magnitude decreasing linearly towards the middle of the ramp. In both cases, the middle read has zero weight, the first half of the reads have negative weights, and the second half have positive weights (the initial reads are subtracted from the latter reads).

Figure 1 shows the optimal weights for each read, normalized to the weight of the final read, for five S/N values. The largest S/N value approaches the photon noise limit where the weights are zero for all but the first and last reads. Intermediate S/N values, corresponding to photon rates that are between the photon noise and the read noise limits, have weights between the values for these two limits.

An MA table will have lower S/N than the case of saving all of the reads because the on-board averaging of reads requires that all reads within a resultant carry the same weight in the ramp fit. In other words, on-board averaging makes it impossible to realize the weights shown in Figure 1. This is more detrimental in some cases than in others. For example, at high S/N, resultants that average the first and/or last read together with other reads will incur the largest S/N penalties. At lower S/N, a resultant symmetrically straddling the middle read (e.g. encompassing reads 8 - 12 in this example) will have zero weight in the ramp fit, and will therefore discard the otherwise useful information in the reads near the

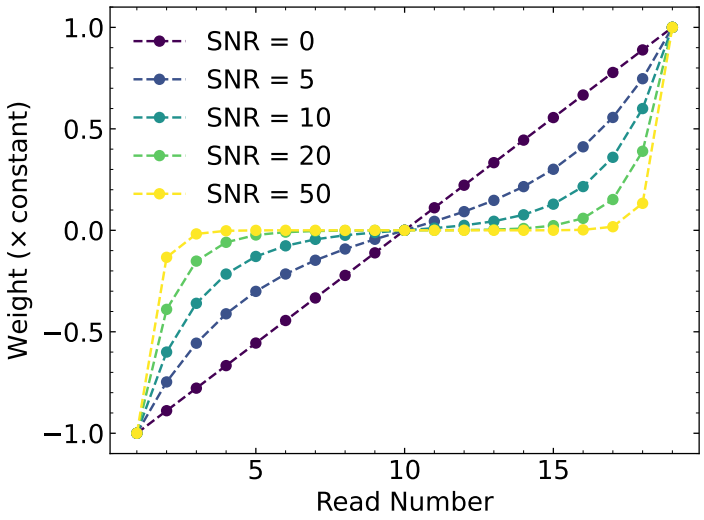


Figure 1: Optimal relative weightings of the individual reads in a 19-read ramp as a function of signal-to-noise ratio of a source. The first and last reads always have the greatest weight, and in the photon noise (high S/N) limit, the first and last reads have all of the weight.

middle.

# 3 A Metric for MA Tables

In the remainder of this technical document, we evaluate MA tables using the ratio of the S/N possible with an MA table to that possible with all of the reads. This ratio will be a function of the count rate. Due to the fact that the weights shown in Figure 1 are no longer attainable, it will always be less than one. Our adopted metric was developed and used in Brandt (2024a); Casertano (2022). Because the signal (the count rate) is the same regardless of the MA tables used, we need only compute the uncertainty on the fit to determine the quality of the MA table. This can be done using the equations of Brandt (2024a).

Both the previous section and Brandt (2024a) only showed results for ramps that do not saturate. We account for saturation here to make results valid for arbitrarily high count rates: all reads with at least 100,000 electrons, and all resultants that reach 100,000 electrons by the end, are treated as unusable. An important advantage of detectors that can be read out nondestructively is their dynamic range; a source that saturates halfway through an exposure can still be recovered. MA tables are intended to preserve this advantage over a CCD-like detector. Cosmic rays manifest as sudden jumps in a pixel's counts. They limit the S/N of a ramp fit, both by making a single read difference unusable and by statistically decoupling the noise before and after the jump. We do not explore sensitivity to cosmic rays in this document, but note that Sharma & Casertano (2024) found that saving the final read without averaging best preserves sensitivity to jumps. This is also desirable for preserving S/N for pixels with high count rates.

The top panel of Figure 2 shows an example of an MA table, while the lower panels show the S/N relative to the case of saving all of the reads. Without saturation (lower left

panel), there is a few percent loss of sensitivity at all count rates. With saturation (lower right panel), there is a small loss of sensitivity if saturation occurs just after the final read of a resultant, a larger loss of sensitivity if saturation occurs later in a resultant, or total loss of sensitivity if saturation occurs before the end of the second resultant. In this case there is no valid resultant difference that can be used to estimate the count rate.

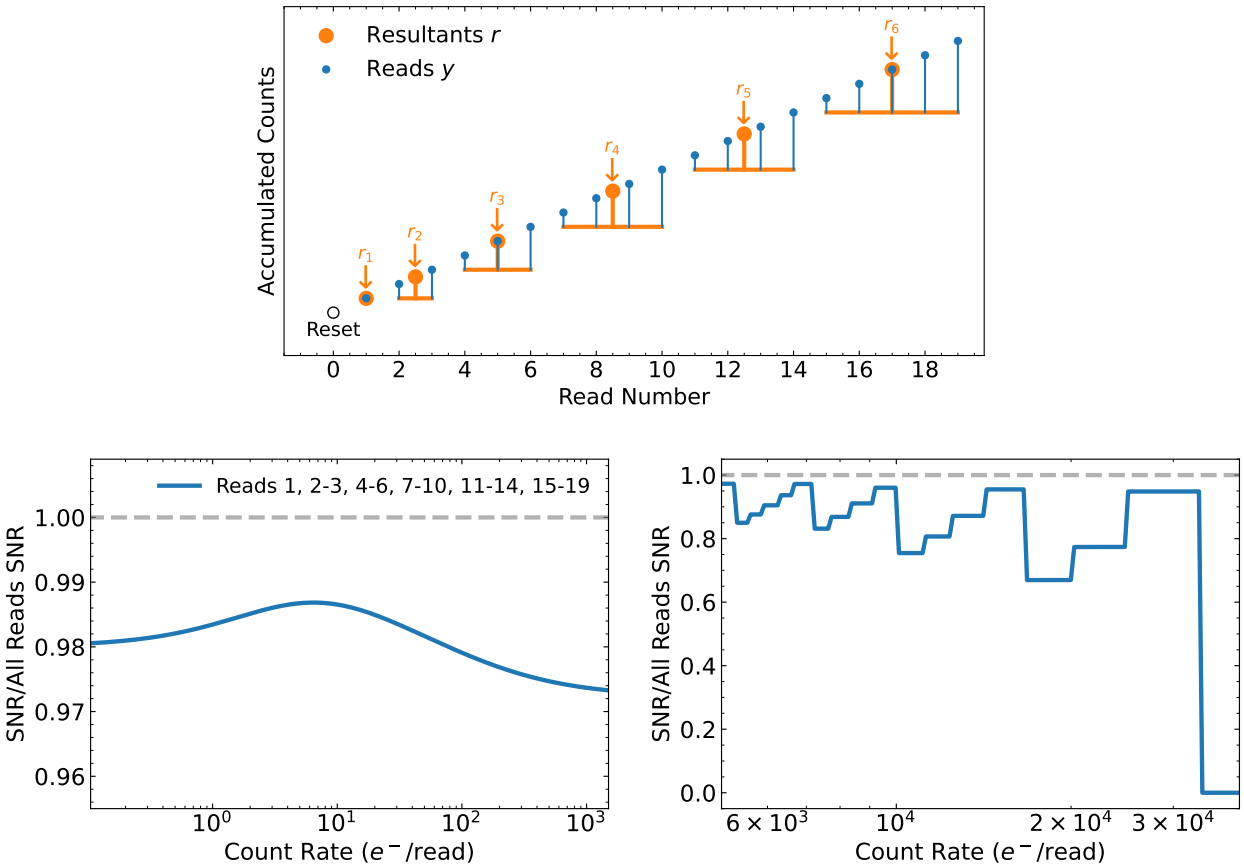


Figure 2: Top: example of combining nineteen reads into six resultants using an MA table. Figure adapted from Brandt (2024a). Bottom: S/N achievable with the MA table shown above relative to the S/N achievable with all of the reads saved separately, assuming a read noise of 10 e$^-$/read. The left panel shows a few percent loss without saturation. The right panel shows much larger losses when saturation occurs partway through a resultant, and total loss when saturation occurs by read 3 (in the second resultant).

The MA table shown in Figure 2 is not optimal. At high count rates, if the final resultant used only one read, we could realize the ideal weights shown in Figure 1. Keeping the final read as its own resultant would also improve sensitivity to cosmic rays (Sharma & Casertano 2024). Omitting the middle (tenth) read entirely from the resultants would likewise offer an improvement in S/N. Improvements in saturation behavior are more subtle, as any MA table will offer tradeoffs here. The following section will present general principles for optimizing an MA table; these will be explored further later in this report.

# 4 General Optimality Principles

Based on the preceding sections, we now list some general principles for constructing optimal MA tables for a given exposure time. These assume that all on-board averaging is unweighted: the reads are simply added together. For optimizing S/N for sources that do not saturate:

- The first and last reads should not be averaged.
- The middle read should not be part of any resultant.
- No resultant should straddle the middle read.

For optimizing sensitivity to cosmic rays:

- The first and last reads should not be averaged.

For optimizing dynamic range, the best approach depends on the distribution of source intensities, but a practical goal might be:

- Fewer reads/resultant should be used earlier in an exposures, and more reads/resultant later in an exposure.

The principles in this section are general guiding ideas that are primarily relevant when the ramp is short and a relatively small number of resultants are used.

# 5 Example MA Table Comparisons

Figure 2 shows an example MA table with 19 reads; the reads are grouped into resultants at 1, 2-3, 4-6, 7-10, 11-14, 15-19. Figure 3 compares this to two alternative MA tables, both of which use the same number of resultants (six) but preserve the final read without averaging and discard reads in the middle of the ramp. The two alternative MA tables offer better performance when the pixel does not saturate, with some tradeoff in S/N when the pixel does saturate during the ramp.

With the final read saved on its own and without increasing the number of resultants, fewer resultants are available for the first part of the ramp. This can result in a loss of S/N for a pixel that saturates during the ramp. The S/N values go to zero at very high count rates where the second resultant is already saturated. Unless read 1 and read 2 are each saved as their own resultants, this will occur at lower count rates than with all of the reads. The apparently large losses in S/N at very high count rates depend on being photon noise limited up to full well ($\sim 10^5$ electrons): detector artifacts and linearity must be corrected to a few tenths of a percent.

Figure 3 does not account for the loss of S/N due to cosmic ray hits. In JWST NIRCam data, a cosmic ray impacts a pixel, on average, once per ≈5,000-10,000 seconds (Brandt 2024b)[1]. The 19-read MA tables shown in Figure 3 take ≈60 seconds to read out, so we

[1] see also https://jwst-docs.stsci.edu/jwst-exposure-time-calculator-overview/jwst-etc-calculations-page-overview/jwst-etc-cosmic-ray-implementationd

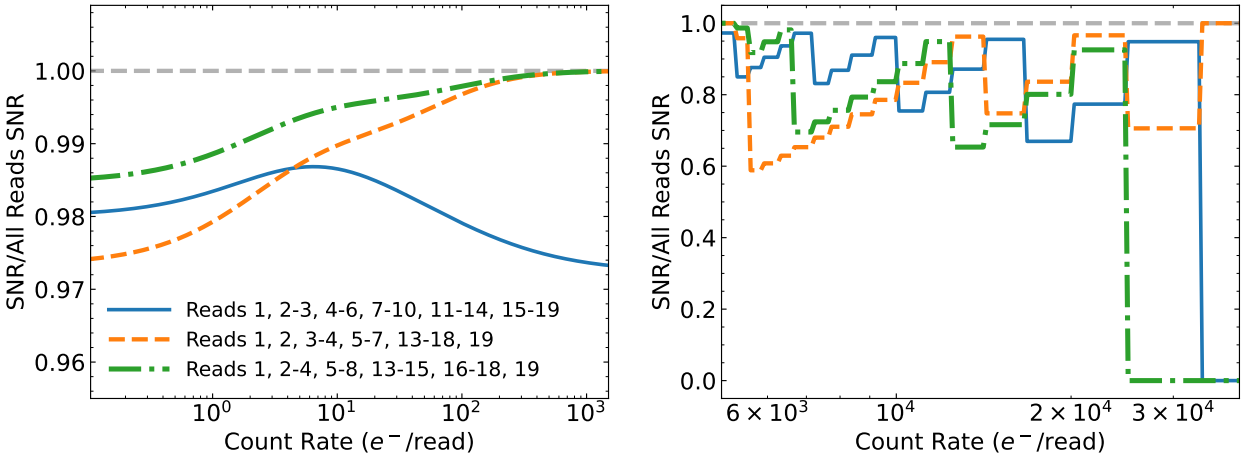


Figure 3: Comparison of three MA tables: the example shown in Figure 2 (blue solid lines), and two other MA tables that drop the middle reads and preserve the final read without averaging. As for Figure 2, note the very different $y$-axis ranges in the left and right panels.

can expect ∼1% of pixels to be affected by a cosmic ray in any given ramp. If a cosmic ray does hit, it will be at a random time within the ramp, and can corrupt either one resultant difference (if it occurs between resultants) or two resultant differences (if it occurs within a resultant).

Figure 4 shows the effect of the MA table choice on S/N loss due to a correctly flagged cosmic ray. The signal-to-noise ratios are calculated as described in  with one or two resultant differences marked as unusable. The shaded regions encompass 68% of the range of losses in sensitivity depending on when in the ramp a cosmic ray hits. The loss of sensitivity is expressed relative to the loss when saving all of the reads. This loss of sensitivity can be large: if a cosmic ray hits in the middle of a long resultant or a period when the reads are being discarded, for example, a large fraction of the integration time may be rendered unusable.

The best MA table for recovering S/N from cosmic rays (as opposed to identifying them in the first place) is one in which the resultants contain equal numbers of reads. Unfortunately, this is strongly opposed to the principles for maximizing dynamic range and S/N for pixels that neither saturate nor receive a cosmic ray. Cosmic rays are expected to affect $\lesssim$1% of 19-read ramps, while the choice of MA table can have a significant effect on the S/N in those cases. As a result, cosmic rays may be thought of as modifying the curves shown in Figure 3 by a couple tenths of a percent.

# 6 A Notebook for Comparing MA Tables

The comparisons between different MA tables shown in this document were made using a Jupyter notebook supplied with this technical report. The notebook uses the `fitramp` module (https://github.com/t-brandt/fitramp) to compute the optimal weighting of the resultants and to compute the uncertainty on the slope. The notebook will make Figures

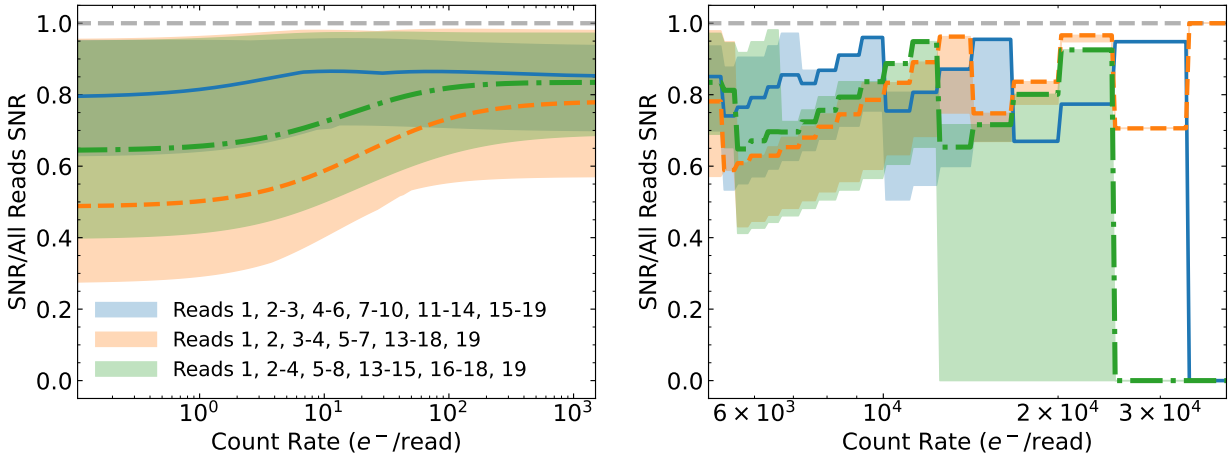


Figure 4: Range (encompassing 68% of possible jump times) of the S/N losses due to a correctly flagged jump for the three MA tables shown in Figure 3. Note the different $y$ axis scale from Figure 3, as losses of half or more of the S/N are now common as a large fraction of the exposure time becomes unusable.

2 and 4 as shown in this document.

A user can compare several MA tables by supplying each as a list of lists. For example, the list of lists

```
readtimes = [[1], [2, 3], [4, 5, 6], [7, 8, 9, 10],
             [11, 12, 13, 14], [15, 16, 17, 18, 19]]
```

would specify the MA table shown in Figure 2. If the user wishes to show the 68% ranges of sensitivity losses due to cosmic rays, they would set the supplied boolean variable to `True`.

# 7 Weighting Reads Within a Resultant

The analysis of Sections 3–5 was restricted to MA tables in which reads are averaged into resultants, with each component read receiving equal weight in a resultant. The Appendix of Brandt (2024a) shows that it is possible to outperform this scenario by carefully choosing weights for each read within a resultant. The improved performance comes from more closely approaching the weights shown in Figure 1.

The suggested approach in Brandt (2024a) is to weight each read within a resultant using the weights that are optimal in the read noise limit. This will enable the user to achieve the same S/N in the read noise limit as if all of the reads were saved. It will also offer improvements in S/N at all count rates (assuming the pixel does not saturate, and that no cosmic ray hits during the exposure). Arbitrary weights can be accounted for in the likelihood-based ramp fitting algorithm, subject only to the constraint that the sum of the weights within a resultant cannot be exactly zero. With weighted resultants, the guidelines in Section 4 to avoid the middle read need not apply. Weighting schemes can also be devised that strictly improve S/N everywhere even if a ramp saturates before the final read. One

example of such a weighting scheme averages all reads with equal weight before the middle of the ramp, and adopts the optimal weights in the read noise limit after the middle of the ramp.

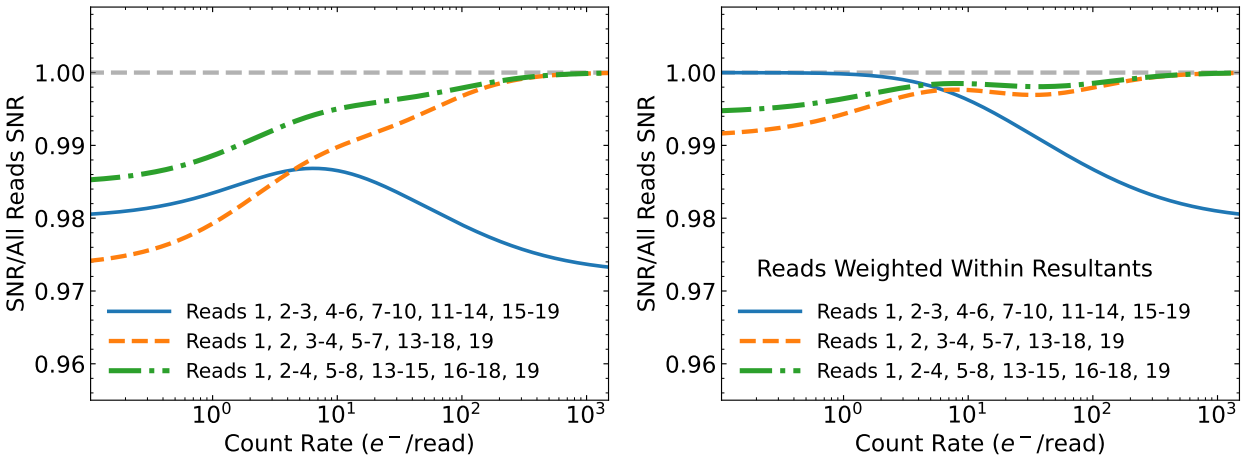


Figure 5: Improvement in S/N possible if reads can be weighted within resultants. Left panel: same as the left panel of Figure 2. Right panel: S/N when the reads are weighted within resultants using the optimal weights in the read noise limit, as in the Appendix of Brandt (2024a).

Figure 5 shows the improvements in S/N that are possible if the MA tables in Figure 2 weight the reads within each resultant. The S/N does not quite reach the S/N with all reads for the two alternate MA tables because some reads are discarded. At all count rates that do not saturate, however, the alternative MA tables that separate out the first and last read both stay within 1% of the S/N achievable with all of the reads.

# 8 Nonuniform Weighting of Reads for Roman

The flight software for Roman will not support a general weighting scheme like that shown in Figure 5. However, some weighting is indeed possible. In this section, we work through one example of a grouping of reads into resultants under three cases:

1. Equal weighting of reads, as discussed in Sections 3–5

2. Unequal weighting of reads with optimal weights in the read noise limit, as suggested by Brandt (2024a)

3. Unequal weighting of reads with weights that can be implemented by the Roman flight software

For the readout pattern in this section, we take 19 total reads, numbered 1 through 19, and grouped as

```
[[1], [2, 3], [4, 5, 6], [7, 8, 9, 10], [11, 12, 13, 14, 15, 16, 17, 18], [19]]
```

The very long penultimate resultant is chosen for several reasons. First, the later in the ramp such a long resultant occurs, the less it impacts dynamic range. Second, the weightings that can be implemented in Roman can make much more effective use of a long resultant after the ramp's midpoint than before it. Finally, the optimal weights are different for a source that saturates as for a source that does not saturate. For the resultants before the ramp's midpoint, our Roman-implementable weights will be equal for all reads. Using the last read as its own resultant is done both for signal-to-noise in the photon-noise limit and to maximize sensitivity to cosmic ray hits that manifest as jumps (Sharma & Casertano 2024).

| Reads | Uniform Weights | Roman-Compatible Weights | Optimal Weights |
|---|---|---|---|
| 1 | 1 | 1 | 1 |
| 2–3 | $\frac{1}{2} \times (1, 1)$ | $\frac{1}{2} \times (1, 1)$ | $\frac{1}{15} \times (8, 7)$ |
| 4–6 | $\frac{1}{3} \times (1, 1, 1)$ | $\frac{1}{3} \times (1, 1, 1)$ | $\frac{1}{15} \times (5, 4, 3)$ |
| 7–10 | $\frac{1}{4} \times (1, 1, 1, 1)$ | $\frac{1}{2} \times (1, 1, 0, 0)$ | $\frac{1}{6} \times (3, 2, 1, 0)$ |
| 11–18 | $\frac{1}{8} \times (1, 1, 1, 1, 1, 1, 1, 1)$ | $\frac{1}{39} \times (1, 2, 4, 4, 4, 8, 8, 8)$ | $\frac{1}{36} \times (1, 2, 3, 4, 5, 6, 7, 8)$ |
| 19 | 1 | 1 | 1 |

Table 1: Weights of the individual reads comprising a resultant for a sample readout pattern. Constants are factored out to more clearly show the relationship between the individual weights.

Table 1 lists three sets of intra-resultant weights for this MA table: uniform weights, the ideal weights in the read noise limit (as suggested by Brandt (2024a)), and weights that are compatible with Roman's flight software. In all cases, constants have been factored out to more clearly show the relationship between the weights of the individual reads. Figure 6 compares the SNR realized by the three intra-resultant weighting schemes listed in Table 1. The Roman-compatible weights realize most of the SNR that is possible from intra-resultant weighting.

The improvements possible with intra-resultant weighting depend on the MA table. Weights close to the ideal read-noise-limited weights can be achieved in Roman for resultants after the ramp midpoint. This is also consistent with the use of MA tables to improve dynamic range: fewer reads per resultant early in a ramp gives a larger number of measurements for a bright source before saturation.

Intra-resultant weighting is possible within the constraints of the Roman flight software, and enables nearly optimal use of long resultants after a ramp's midpoint. The use of fewer resultants could be important in limiting the downlink bandwidth requirements of Roman without incurring a SNR penalty.

# 9 Conclusions

This technical document summarizes the use of MA tables by the Roman Space Telescope and provides a quantitative metric to compare different MA tables. It also shows that it is possible to further improve the performance of the Roman mission by weighting the

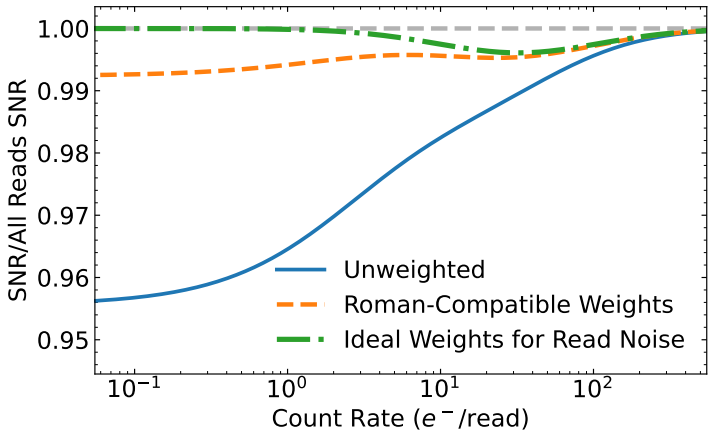


Figure 6: Comparison of the three intra-resultant weighting schemes listed in Table 1. For this MA table, with six resultants comprising reads 1, 2-3, 4-6, 7-10, 11-18, and 19, the Roman-compatible weighting scheme achieves most of the SNR gain that is possible with intra-resultant weighting. Read noise is assumed to be 10 $e^-$/read.

individual reads that combine to form resultants in an MA table. A companion Jupyter notebook enables a user to perform comparisons themselves between any pair of MA tables.

An MA table reflects a compromise between saving all reads individually, which offers the best data quality and the most resilience to cosmic rays, and averaging reads into resultants, which requires a much lower downlink bandwidth. This document provides a tool to help choose the best MA table for a given science case. We also recommend that the mission study the possibility of weighting individual reads that comprise resultants, which could enable the preservation of SNR for a lower downlink bandwidth than would otherwise be possible.

# 10 References

Brandt, T. D. 2024a, PASP, 136, 045004

—. 2024b, PASP, 136, 045005

Casertano, S. 2022, Determining the Best-Fitting Slope and Its Uncertainity for Up-the-Ramp Sampled Images with Unevenly Distributed Resultants, Technical Report Roman-STScI-000394, Baltimore: STScI

Sharma, S., & Casertano, S. 2024, PASP, 136, 054504