Balancing Adjustment Study,
Flexible Flow Management Program 2017

Office of the Delaware River Master, 2024

DRAFT

Table of Contents

Introduction

An empirical model, incorporating a recession and routing strategy, is used by the Office of the Delaware River Master (ODRM) for forecasting the average daily flow at Montague to determine whether directed releases are required from New York City (NYC) reservoirs to maintain the flow objective at Montague. These procedures have been refined over nearly six decades of use and generally give good results during periods of relatively stable flow.

However, as stated in the 1958 ODRM annual report, "*It is impossible, because of the many factors involved, to precisely predict in advance the resultant flow at Montague. However, over a period of time any cumulative error in the estimating procedure should be eliminated as far as possible. To accomplish this, a balancing adjustment is applied to the indicated deficiency or excess in computing the directed release. This adjustment is based on the amount by which the directed release is greater or less than the release actually required to maintain the prescribed rate of flow at Montague….*" This was the first instance that a balancing adjustment was described in an annual report and ODRM has utilized different procedures to calculate the amount since that time. The use of the balancing adjustment is a means to conserve the waters of the basin by tracking and compensating for the inherent errors in forecasted directed releases.

As part of the Agreement for a Flexible Flow Management Program 2017 (FFMP 2017), it was determined that the "Decree Parties shall study, evaluate, and consider the River Master's balancing adjustment procedure." In collaboration with USGS technical experts and the Decree Party Work Group, ODRM prepared and executed a study to accomplish this task.

Purpose and Scope

The purpose of this report is to document the process, findings, feedback, and final recommendations for improving the ODRM balancing adjustment procedure. The Statement of Work for the study is in the Appendix of this report, and available at (https://webapps.usgs.gov/odrm/documents/Balancing-Adjustment-Study-SOW-20210413.docx). The tasks of the study included the compilation of data, development of a model to replicate the directed release design process, an assessment of the balancing adjustment, and development and testing of alternative procedures. A description of the balancing adjustment procedure, model structure, data assessment and results are presented, as well as the feedback from the Decree Party Workgroup which guided the recommended modifications.

Directed Releases and the Balancing Adjustment Procedure

ODRM directs releases from NYC reservoirs as needed to maintain flow targets on the Delaware River at Montague, NJ (Montague). The directed releases are determined 3 days in advance to allow for time of travel from the reservoirs to Montague. The forecasted flow components include uncontrolled runoff, baseflow, and powerplant releases, each with its own inherent error. Examples of the types of error are shown in Table 1.

Table 1. Sources of directed release forecast error.

Flow component	Examples of error
Powerplants	• Forecast changes after directed release design calculation • Power market fluctuations • Maintenance emergencies
Runoff from rainfall	• Changes in weather • Accuracy of prediction tools (for example, often "tuned" to high flows) • Vegetation growth
Baseflow behavior	• Changing groundwater-water interactions
Streamflow	• Shifts in discharge ratings due to channel vegetation

After the forecast date has passed, the amount of flow that was actually required to meet the flow target is determined using observed streamgage and release data, assuming a "perfect forecast". A *balancing adjustment* can be calculated to account for the differences between the directed and the actual required releases. When directed releases are determined to have been excessive (greater than required), then the balancing adjustment reduces the subsequent directed release requirement. Conversely, if prior directed releases have been insufficient, then the balancing adjustment increases the directed release requirement. The balancing adjustment is applied to the following day's directed release based on the previous day's cumulative error. Currently, the determination of the balancing adjustment requires calculating the difference between the cumulative actual directed release and the cumulative directed release required with perfect forecasting, with both quantities accumulated beginning June 15th of each year. Positive cumulative error indicates the river "owes" the reservoirs; negative error indicates the reservoirs "owe" the river. The balancing adjustment is currently computed as 10 percent of the cumulative error. The value is then multiplied by -1 so that the applied adjustment "balances" the directed release in the proper direction:

$$\boldsymbol{Balancing\ Adjustment} = \left[\frac{(\boldsymbol{CDIR - CARR})}{\boldsymbol{10}}\right] \times -\boldsymbol{1} \qquad \textbf{Eq. 1}$$

With Balancing Adjustment limited such that
-50 cubic feet per second (ft^3/s) < Balancing Adjustment < 50 ft^3/s

Where:

CDIR = Cumulative Directed Releases in response to ODRM (ft^3/s-days),

CARR = Cumulative Actual Required Releases (ft^3/s-days), where:

Actual Required Release (ARR)=

Flow Objective – (Previous day's daily mean discharge – NYC Releases) (ft^3/s).*

**Not including Interim Excess Quantity Releases*

Using a percentage of the cumulative error is intended to distribute the balancing adjustment, moving the balance toward zero over time. The maximum daily balancing adjustment is limited to +/- 50 ft^3/s to preclude unacceptably large variations in the adjusted flow objective. The cumulative error is reset to zero each year on June 15. See a simplified example of the data and computations of the balancing adjustment in Table 2 Table 2 and Figure 1. **Error! Reference source not found.**In the example, based on the observed data (from yesterday) on June 15th, a balancing adjustment of -15 ft^3/s would be applied to the directed release forecast for June 19th.

Date	Computed deficiency from flow target	Directed release in response to ODRM	Actual required release	Cumulative directed	Cumulative required	Cumulative error (cumulative directed minus cumulative required)	Balancing Adjustment
June 15	550	550	400	550	400	150	-15
June 16	600	600	450	1150	850	300	-30
June 17	650	650	500	1800	1350	450	-45
June 18	660	660	550	2460	1900	560	-50
June 19	340	325	360	2785	2260	525	-50
June 20	490	460	500	3245	2760	485	-49
June 21	620	575	650	3820	3410	410	-41
June 22	650	600	680	4420	4090	324	-33

Table 2. Simplified example of balancing adjustment calculation where the cumulative error is positive. The balancing adjustment is equal to 10% x -1 of the cumulative error.

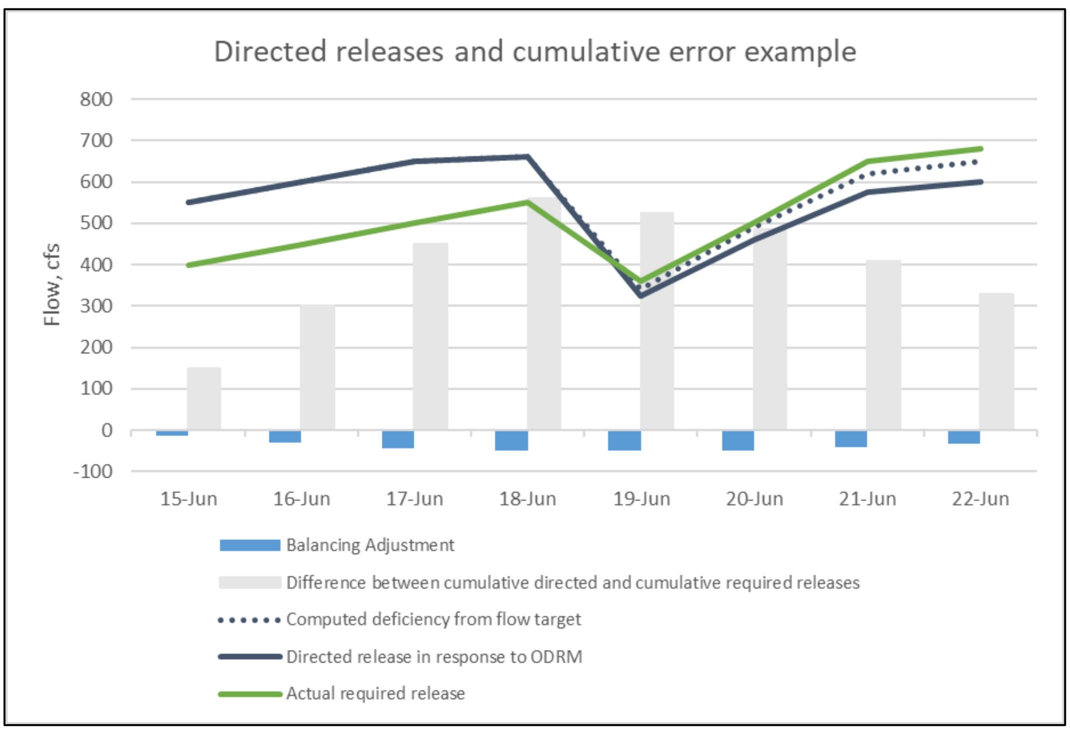

Figure 1. Plot showing an example of directed releases, actual required releases, cumulative error, and the balancing adjustment.

Some pre-existing criticisms include that the process was overly complex, isn't effective, and carries a balance for long periods of time. A cursory review of the balancing adjustment procedure in 2018 immediately resulted in a proposed refinement by ODRM. Conservation releases from NYC reservoirs (releases made regardless of whether a directed release is required) were essentially not yet established or very minimal when the procedure was originally implemented in 1958. During the 1970s, an experimental program of enhanced releases for fisheries became part of the management plan for the Delaware River through unanimous agreement by the Decree Parties to DRBC docket D-77-20. Though the volumes have evolved through multiple D-77-20 revisions and FFMPs, such "conservation releases" have been included in the management plans in use since that time. The balancing adjustment continued to be calculated on forecast directed releases versus actual required releases, without consideration of effects of conservation releases.

Occasionally, conservation releases can be greater than releases directed to meet the flow target at Montague. In these instances, the forecasting error does not impact the operations of the reservoirs and the flow target is met, yet the error is still accumulated and applied in the balancing adjustment. The proposed refinement to the procedure was to compute the cumulative error and apply the adjustment only when directed releases were greater than the conservation releases. The change was implemented in 2018, with a more detailed assessment planned for the following years as part of the FFMP 2017 Balancing Adjustment Study.

Past Performance

Directed release and balancing adjustment data from 2010-2019 were available for assessment at the time of this study. The directed release error on a daily basis ranged from -1,162 to 1,350 ft^3/s, with a median value of 25 ft^3/s (Figure 2 and Table 3. Summary statistics of directed versus actual required releases from NYC reservoirs, 2010-2019.). The largest errors typically occurred due to imprecise forecasting in the timing and/or magnitude of summer thunderstorms. Based on the interquartile range, directed releases were generally within +/- 200 ft^3/s of actual required flow. In examining the cumulative error, directed releases were more often greater than actually required. With the cumulative error values skewed in the positive direction, the cumulative range was from -4,103 to 9,219 ft^3/s, with a median value of 538 ft^3/s (Figure 3 and Table 3Table 3). The interquartile range was from -434 to 4,767 ft^3/s. Most years ended with a large positive cumulative error and an unresolved negative balancing adjustment. Exceptions to this typical condition occurred in 2018 and 2019. The inclusion of conservation release criteria and an adjustment of the baseflow recession curve, both implemented in 2018, may have contributed to this improvement. There were also fewer directed releases these years compared to years past.

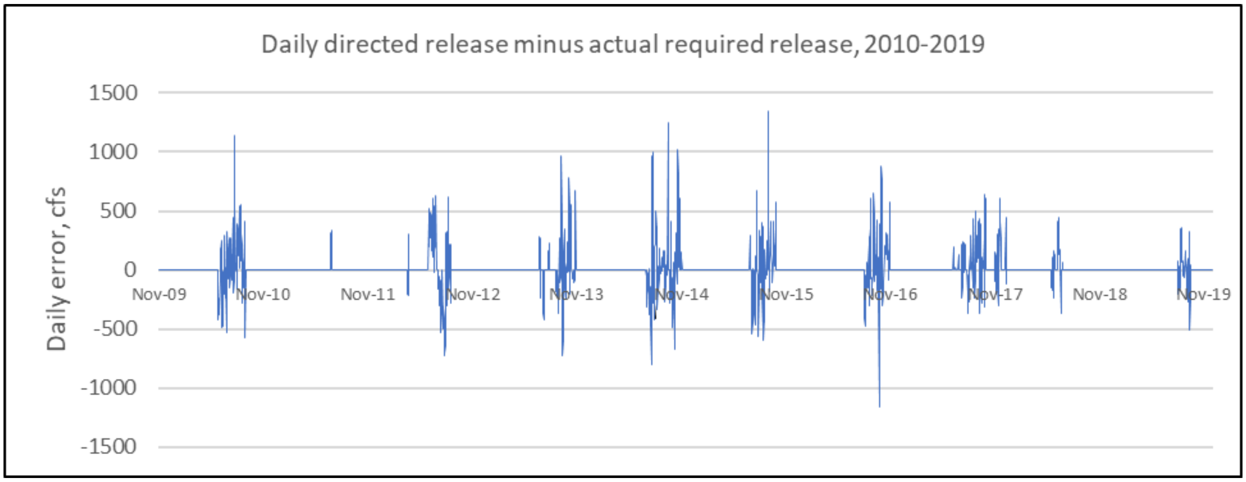

Figure 2. Directed release minus actual required release to meet flow target at Delaware River at Montague, NJ, 2010-2019.

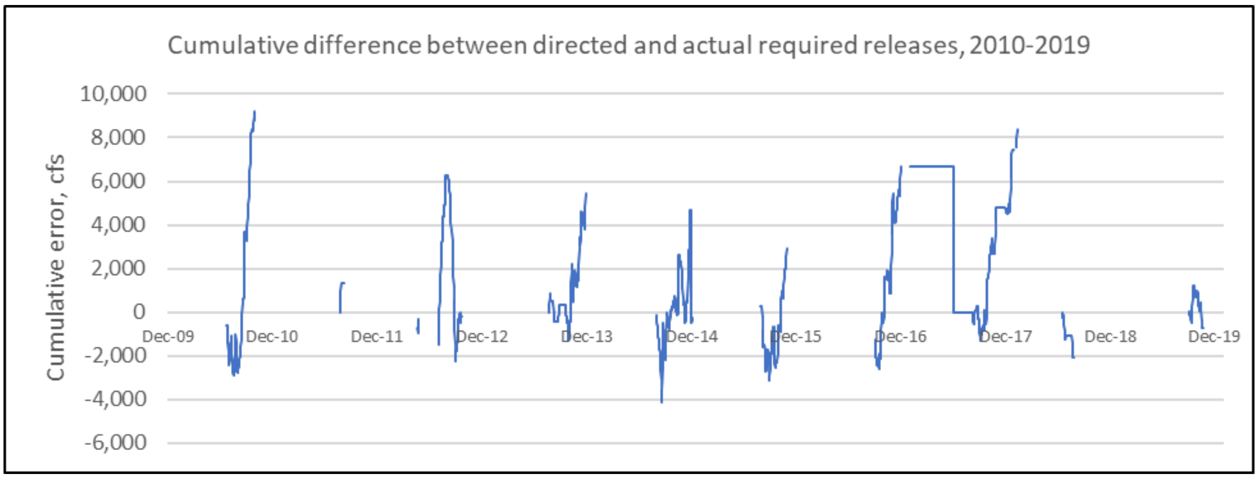

Figure 3. Cumulative error of directed release with respect to actual required release to meet flow target at Delaware River at Montague, NJ, 2010-2019.

Table 3. Summary statistics of directed versus actual required releases from NYC reservoirs, 2010-2019.

Directed release versus actual required, 2010-2019		
	Daily error	**Cumulative error**
Minimum	-1,162	-4,103
10th percentile	-326	-1,728
25th percentile	-166	-434
Median	10	538
75th percentile	197	4,767
90th percentile	396	6,697
Maximum	1,350	9,219

Proposed Improvements

A virtual workshop was held July 1, 2021, comprised of the Balancing Adjustment Study Team, USGS experts, and Decree Party Workgroup members. Available data and past performance of the balancing adjustment were reviewed, and a brainstorming session was held which produced more than 40 individual suggestions. The list of possible improvements was culled and distilled into five categories.

Components review – Examine when and how to accumulate error and apply the balancing adjustment in relation to conservation releases. Review the different components of the algorithm to ensure accurate accounting.

Reset – Determine impacts of using a different reset date than June 15th, including not resetting to zero or resetting when reservoirs spill.

Distribution – Evaluate the use of a percentage of the cumulative error and consider revising from the current value of 10% (corresponding to the value of 10 in the denominator of Eq. 1).

Caps – Explore the effects of different cap values to expedite the resolution of the cumulative error. Suggestions included the use of a variable cap and no cap. The current cap is 50 ft^3/s.

Removal – Consider the removal of the balancing adjustment and only track the forecast error.

Compile history – Review past ODRM reports to build a history of the Balancing Adjustment and expedite the transcription and addition of older ODRM data into the Aquarius database to explore effects of balancing adjustment during drought (e.g., 1960s, 2001-02).

Modeling the ODRM Directed Release Design Process and Evaluating Alternate Approaches to the Balancing Adjustment

To facilitate a detailed assessment of the historical performance of directed release forecasts and associated error, and to evaluate alternate approaches of the balancing adjustment, a Python model reproducing the design process, herein referred to as "pyDR," was developed. The schematic of this model is given in Figure 4. A list of frequently used abbreviations to describe the model flow components are provided here for ease of reference.

ARR	Actual required release
BA	Balancing adjustment
CR	Conservation release
DR	Directed release
IERQ	Releases from Interim Excess Release Quantity banks (not directed by ODRM)
Q_{NYC}	Total release from Pepacton, Cannonsville, and Neversink Reservoirs

Model Description

On a given day, the first set of calculations in pyDR (Figure 4, upper box) takes as input the previous day's values for

- combined release from New York City (Q_{NYC}) reservoirs at Pepacton, Cannonsville, and Neversink, lagged by their respective travel times to Montague and summed,
- preliminary observed flow at Montague,
- lagged reservoir releases from two reservoirs operated primarily for power generation, Wallenpaupack and Rio, and
- the ODRM directed release.

These data are used to estimate the non-reservoir component of the previous day's Montague flow and the Balancing Adjustment (BA). One additional input used in some applications of the model is the previous day's total conservation release (CR) from the NYC reservoirs. The approach starting in 2018 was that if this value was greater than the previous day's directed release (DR) or if DR was zero, then the BA calculations were not carried out. Specifically, when CR>DR or DR=0, the cumulative sums that are the penultimate step in the BA calculations were not accumulated, and the BA was set to zero. For most of the historical record considered in the present analysis (2010-2017), however, this criterion was not applied. Therefore, the pyDR runs intended to reproduce the historical period do not include this criterion prior to 2018. In the runs evaluating alternate DR or BA calculation approaches, the criteria constituting a particular scenario are applied for the entire study period.

The next set of calculations in pyDR (Figure 4, lower box) uses the BA and, in practice, the non-reservoir Montague flow component, along with values forecast for 3 days after the current day, to calculate the DR. In practice, baseflow and prior runoff recession estimates of the non-reservoir Montague flow component were often adjusted by ODRM using professional judgment. Because human judgment had been applied to produce this forecast and this judgment is not easily captured algorithmically, the historical values for the resulting forecast are used in the model, rather than using any function of the non-reservoir Montague flow component. The forecast values that are used in the DR calculations are 1) the forecast total rainfall runoff, which is calculated outside of the model using a method developed by ODRM and is based on the forecast precipitation for the contributing areas above Montague, provided by NOAA; and 2) the forecast releases for Wallenpaupack and Rio reservoirs (lagged by average travel time to Montague), provided by their operators. Using these inputs, the DR needed to meet the target flow in 3 days (the target flow is another model input) is estimated and adjusted using the BA value calculated in the first step.

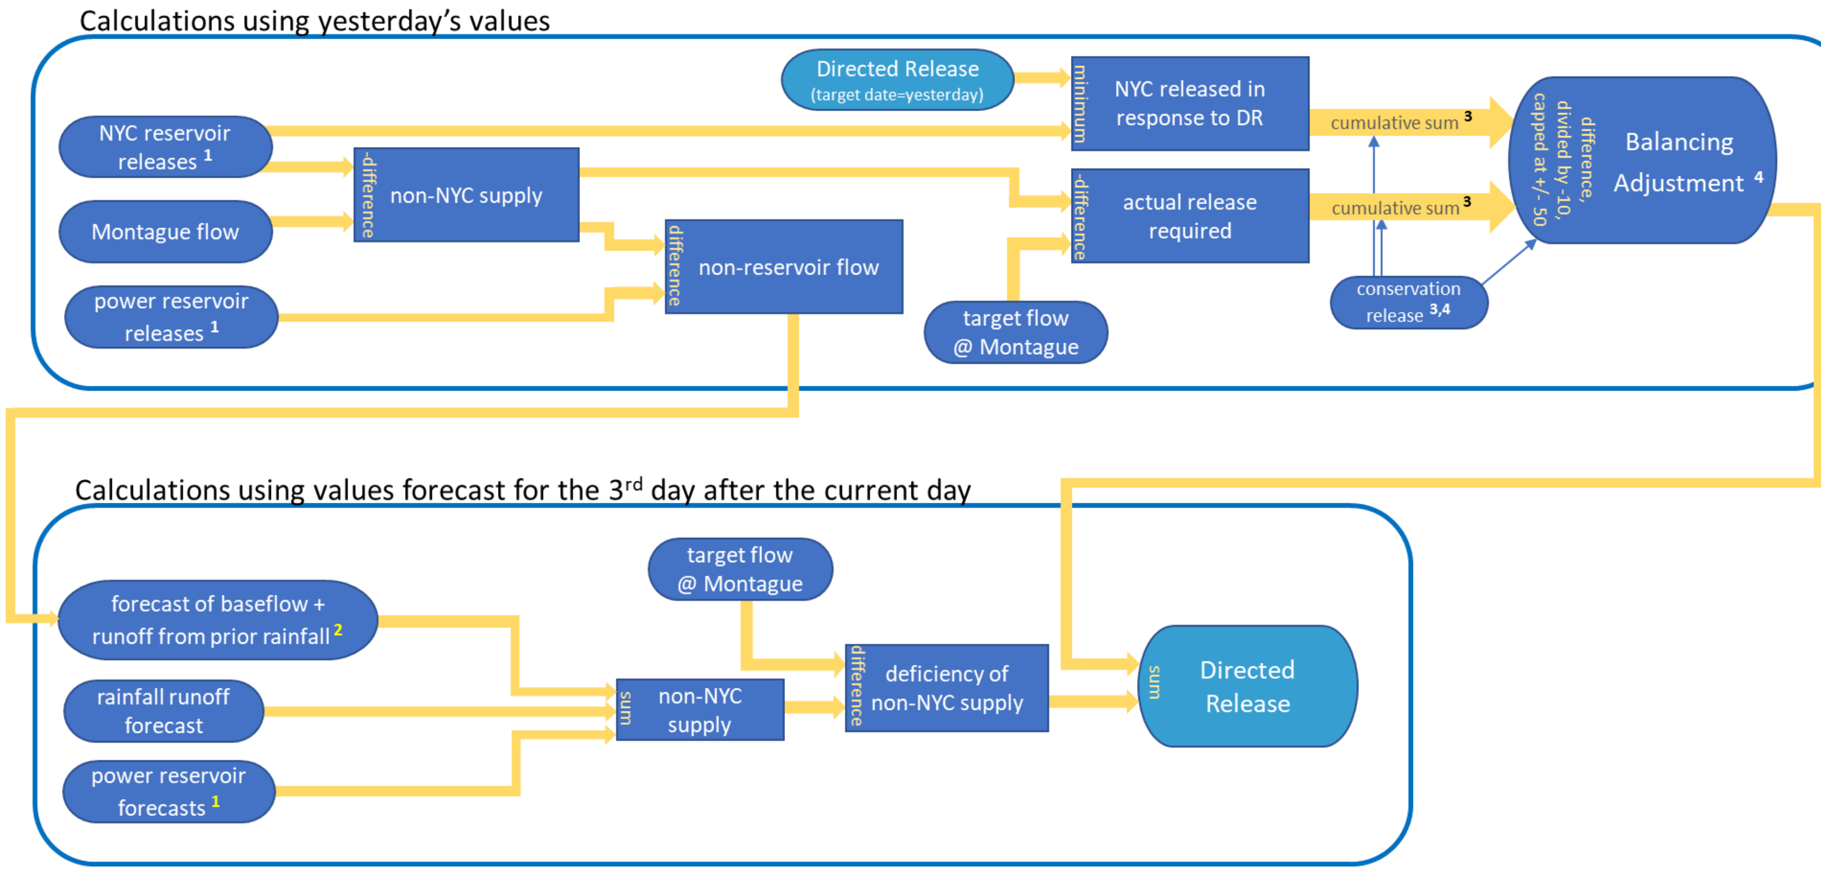

Figure 4. Schematic for pyDR, a model of ODRM directed release. Quantities in the upper box are based on data collected on the day before the current day. Quantities in the lower box are forecasts for the "design date," 3 days after the current day. Differences are top input minus bottom input unless the reverse is indicated by "-difference".

[1] Reservoir releases are lagged by average travel time to Montague.

[2] Historical present conditions forecast ("baseflow + runoff from prior rainfall") values are used rather than calculated values because human judgment was involved as opposed to a strict algorithm.

[3] Cumulative sums are reset to 0 on June 15 and, for 2018-2019, are updated only when yesterday's directed release was greater than 0 and the conservation release was greater than the directed release. This is the significance of the thin arrows from the conservation release input to the cumulative sum arrows; it is not included in the cumulative sums but determines when the sums are accumulated.

[4] The balancing adjustment is 0 unless yesterday's directed release was greater than 0 and the conservation release was greater than the directed release, in which case it is calculated as shown. This is the significance of the thin arrow from the conservation release input to the balancing adjustment box.

Data and Assumptions

All data used were assembled and saved by ODRM in the course of their historical operations, with the exception of NYC conservation release data, which was provided to ODRM by the NYC Department of Environmental Protection (DEP), along with updated NYC reservoir releases, in October of 2021 for use in this study. The data used cover the period 12/01/2009 through 12/31/2019 with the dates being the "design date." The calculated DR value is designed in advance of the design date with the goal of meeting the target flow on the design date.

One challenge with the historical ODRM data was that there are times when the required forecast data were missing because no directing was occurring, and other times when forecast data were missing but should have values of zero. It was very difficult to distinguish consistently between these two cases using other quantities involved in the calculations due to the occasional practice of entering estimated data that would intentionally result in a DR of zero at times when flows were so high that it was known that DR should be zero. This practice no longer occurs. The approach used to address these cases was to assume that DR=0 any time that the forecast of baseflow plus runoff from prior rainfall was missing (inspection of the data showed this to be a good indicator that directing did not occur), or when the flow at Montague was greater than 3,500 ft^3/s (historically, 98.5% of days with non-zero DR occurred when the flow at Montague was less than 3,500).

Another assumption built into the model is that the value for "NYC released in response to DR" is the minimum of DR and the total release from the NYC reservoirs, Q_{NYC} (see Figure 4). When pyDR was first being developed, values for this quantity were not available, so this formula was used instead. Values for "NYC released in response to DR" have since been provided by the NYC DEP and agree well with the values calculated using the formula (R^2=0.96). This formula continued to be used in pyDR instead of the historical data because the formula allows a realistic simulation of alternate scenarios in which both the DR and the NYC releases may change.

There are also individual days and longer periods in the ODRM record in which the BA and/or the DR calculations were carried out differently than the procedure used in pyDR and shown in Figure 4. These appear to be primarily due to a combination of occasional human error and intentional changes in procedure, although in some cases, differences between the NYC reservoir release data used historically and the updated dataset provided by the NYC DEP for this study result in differences between historically calculated values and those calculated by pyDR. In using the Python model to simulate the 2010-2019 period to provide a baseline for further evaluations, it is assumed that the pyDR algorithm described above and depicted in Figure 4 applied throughout the study period.

IERQ releases have specified purposes other than the Montague flow target and are excluded when determining the actual required release at Montague. The amounts, uses, and definitions of IERQ releases evolved over time, and were not always excluded when directing releases to meet the flow target at Montague. For this study, the model excluded the IERQ releases in determining the actual required release for the entire period.

Model Performance

The model does an excellent job of reproducing DR values for the historical period 12/01/2009-12/31/2019 (R^2>0.99, Figure 5a). Due to many instances in the historical record of deviations from the nominal BA calculation procedure, the model's reproduction of the daily change in the difference of the cumulative sums used to calculate BA (Figure 4, upper box) is not as impressive, though still good (R^2=0.86). Because these deviations accumulate in the cumulative sums between the June 15th reset dates, the resulting simulated cumulative sums and subsequently calculated balancing adjustments are in poorer agreement with the historical values (R^2=0.68 and R^2=0.49 respectively, not shown). Such disagreements are the result of model assumptions not always agreeing with the historical data (e.g., in the case of the assumed behavior of "NYC released in response to DR" discussed above) and occurrences in the historical record of the nominal ODRM procedure not being followed.

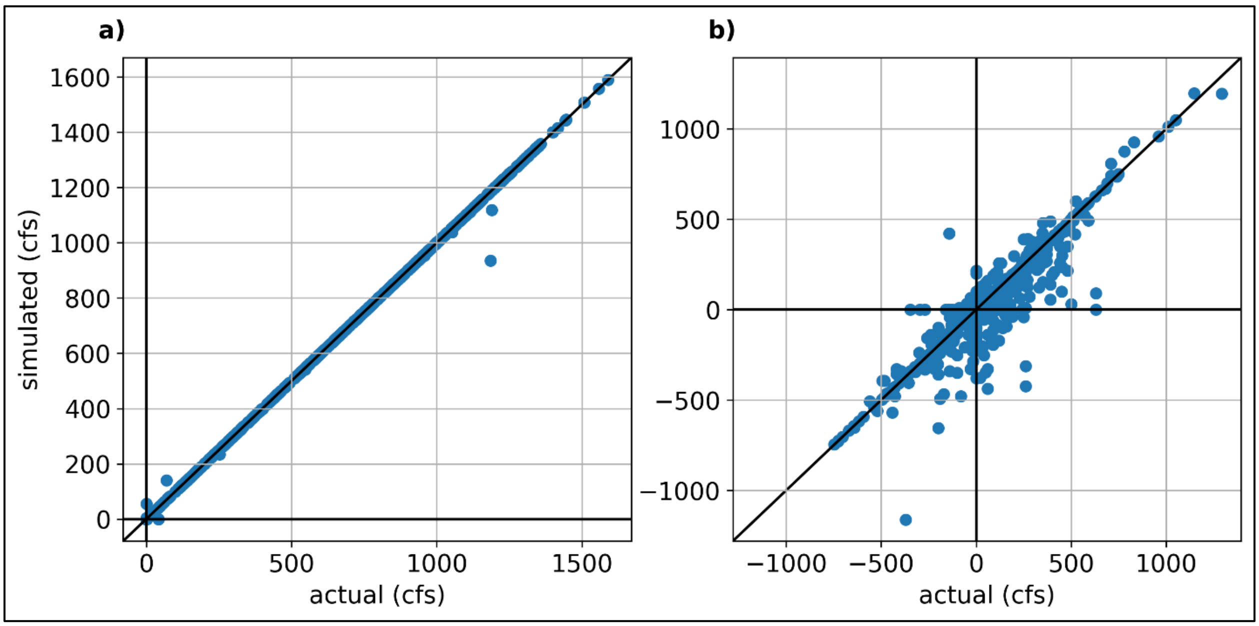

*Figure 5. pyDR performance. Simulated versus observed daily values for 12/31/2009 through 12/31/2019 for **a)** directed release without balancing adjustments applied, and **b)** daily change in difference of cumulative sums used to calculate the balancing adjustment (see Figure 1).*

Evaluation of Alternative Approaches to the Calculation of the Balancing Adjustment

The intent of this study is to evaluate multiple alternative approaches to BA calculation in terms of their effect on Montague flows and on the cumulative error carried through time. This cumulative error is the difference of the cumulative sums (Figure 4, top box), and represents the concept of the NYC reservoirs "owing" water to the Delaware River (negative cumulative error) or the reverse (positive cumulative error). Maintaining a cumulative error as close to zero as possible is one goal in the effort to improve the BA calculations. At the same time, a BA implementation should not cause deviations of Montague flows from the target flow to become unacceptably large.

Based on the discussions in the July 2021 workshop on possible improvements to the balancing adjustment procedure, several alternative scenarios were developed to be run within the pyDR model (Table 4). The alternatives modeled included different cap values, removal of the 10% distribution (by changing the value of 10 in the denominator of Eq. 1 to the value 1), and different reset dates and conditions.

Alternative name	**Cap (ft^3/s)**	**Distribution**	**Reset**
Baseline	50	10	June 15
No balancing adjustment	NA	NA	NA
100 ft^3/s cap	100	1	1) Jan 1 2) reset at spill and > 5,000 ft^3/s
200 ft^3/s cap	200	1	1) Jan 1 2) reset at spill and > 5,000 ft^3/s
400 ft^3/s cap	400	1	1) Jan 1 2) reset at spill and > 5,000 ft^3/s
Variable cap	Dependent on cumulative error	1	1) Jan 1 2) Reset at spill and > 5,000 ft^3/s

Table 4. List of alternative scenarios modeled in the pursuit of an improved balancing adjustment procedure.

Adjustments to Data and Model

Alternate approaches to BA calculation alter the BA values, which would in turn affect NYC reservoir releases and Montague flows, which are inputs to pyDR. Therefore, to create more realistic evaluations, the historical BA components of Q_{NYC} and of the resulting flow at Montague are estimated and removed from those time series in advance of the model run. Then, at each of the model's daily time steps, the newly calculated BA value is incorporated back into those time series. This allows a more accurate depiction of the effects of BA on Q_{NYC} and the resulting Montague flows than using the unaltered input time series.

The details of how the BA component of Q_{NYC} and Montague flows are calculated are shown in Figure 6.

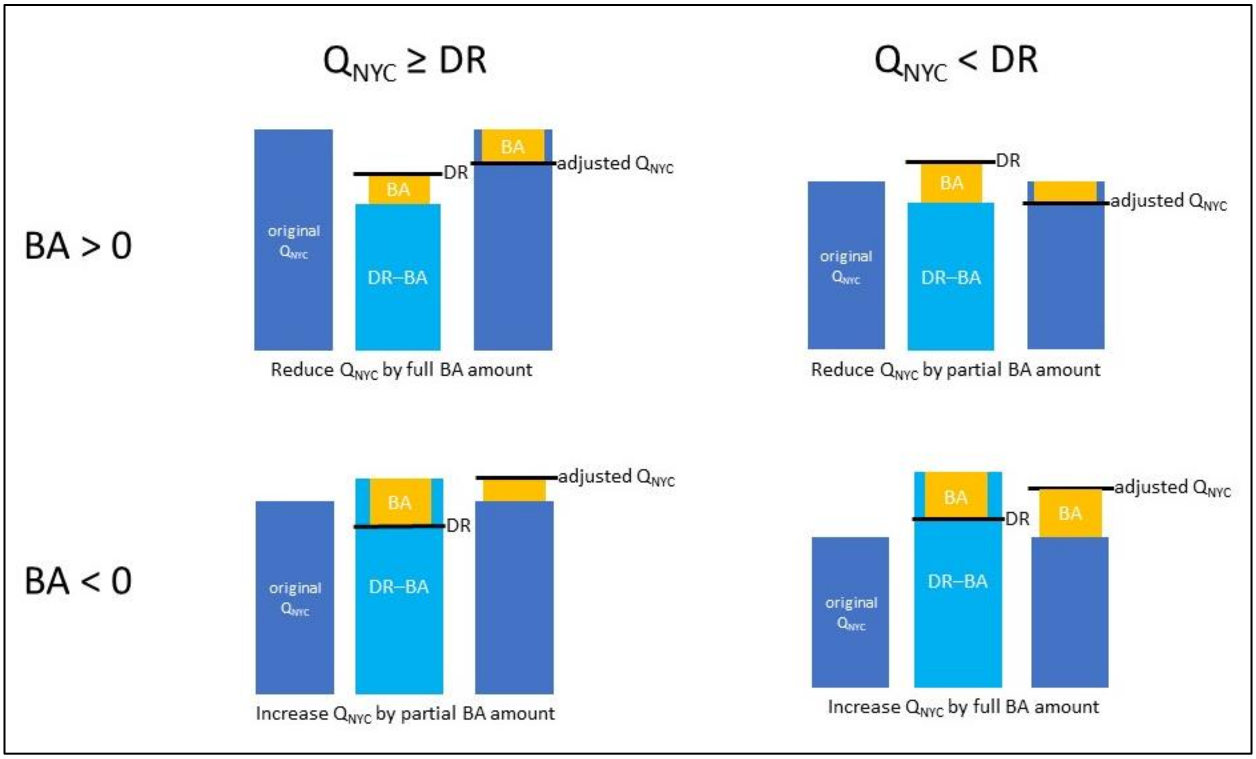

Figure 6. Logic for removing the historical balancing adjustment signal from New York City reservoir releases and Montague flows. BA=balancing adjustment, Q_{NYC}=total New York City reservoir release, DR=directed release. The quadrants represent conditional scenarios based on the value of BA relative to zero (rows) and the relative values of Q_{NYC} and DR (columns). In each quadrant, dark blue bars represent Q_{NYC}, light blue is DR minus BA, and gold represents the BA. The left bar is the historical Q_{NYC} for a given day, the middle bars show the components of DR, and the right bars show the BA-adjusted Q_{NYC}. BA=balancing adjustment, Q_{NYC}=total New York City reservoir release, DR=directed release.

In the upper-left (Q_{NYC}≥DR and BA>0) and lower-right (Q_{NYC}<DR and BA<0) quadrants of Figure 6, the relative values of BA, Q_{NYC} and DR are such that the full BA may be removed from Q_{NYC}. A negative BA is added to Q_{NYC}, and a positive BA is subtracted, to get the adjusted Q_{NYC}. For the lower-left (Q_{NYC}≥DR and BA<0) and upper-right (Q_{NYC}<DR and BA>0) quadrants, the relative values allow only a portion of the BA to be removed from Q_{NYC}, such that the adjusted Q_{NYC} = DR-BA. For the lower-left scenario, the underlying assumption is that the portion of Q_{NYC} in excess of DR reflects an unmet portion of the negative balancing adjustment. Therefore, that portion does not need to be added back in when adjusting Q_{NYC}. For the upper-right scenario, the assumption is that the portion of Q_{NYC} in deficit of DR reflects an unmet portion of the positive balancing adjustment, and that portion does not need to be subtracted when adjusting Q_{NYC}. The one exception to this scheme is for the cases with Q_{NYC}>DR (left column) when CR>DR. In those cases, no adjustment to Q_{NYC} is made because the conservation release overrides the DR. The same resulting adjustment that is made to Q_{NYC} on a given day (remembering that Q_{NYC} represents reservoir releases that are lagged by their approximate travel time to Montague) is made to Montague flows, thus removing the estimated BA influence from both time series.

As the pyDR model runs, the calculated BA portion of DR for each design date is used to reintroduce the BA signal into that day's Q_{NYC} that and the flow at Montague by adding it to both values, limiting the adjustment such the resulting Q_{NYC} is not less than CR.

One final adjustment to the input data was needed to compensate for discrepancies between the updated NYC DEP Q_{NYC} and CR data and the historical Interim Excess Release Quantity (IERQ) data. The IERQ releases have specified purposes other than the Montague flow target and are therefore excluded. There were a limited number of days during the study period when Q_{NYC} was equal to CR but IERQ was greater than zero. Q_{NYC} minus IERQ should not be less than CR, so IERQ was adjusted on those days so that:

$$IERQ_{adjusted} = \text{minimum}\ (IERQ_{original}, Q_{NYC} - CR).$$

Results and Feedback

The model evaluations of the proposed improvements were reviewed by the Study Team, USGS experts, and the Decree Party Workgroup in at a second virtual workshop on July 20, 2022.

Components review – During the initial workshop it was recognized that in addition to varying the parameters of the core BA calculation algorithm, the algorithm itself may benefit from added scrutiny, especially the modified procedure that had only begun in 2018. The procedure took conservation releases into account, and its use had not been assessed to date. A detailed examination of the balancing adjustment core algorithm was undertaken in the course of modeling the various alternatives. The model provided an opportunity to ensure that the algorithm was operating as intended, providing the most fair and logical representation of the cumulative error within the context of conservation releases. Two conditions and one quantity were identified as the core components of the BA calculations:

1) when to apply the resulting balancing adjustment value to the directed release
2) when to accumulate the error, and
3) what to accumulate for the error.

In the detailed examination of the balancing adjustment calculations, several potential changes were considered, and a proposed new approach for determining these three factors was arrived at. Table 5 shows the current[1] and proposed approaches. The components used in the approaches are abbreviated as follows:

DR	Directed Release
CR	Conservation Release
Q_{NYC}	Total release from Pepacton, Cannonsville, and Neversink Reservoirs
ARR	Actual Required Release
min	minimum
max	maximum
IERQ	Interim Excess Release Quantity

[1] The term "current" is used to refer to the ODRM operating procedure in effect at the time this work was carried out. That procedure was in effect from July 2018 through June 2023.

Table 5. Current and proposed base approaches to calculating the balancing adjustment.

	When to Apply	When to Accumulate	What to Accumulate
Current Approach	when directing and DR>CR	when directing and DR>CR	min(Q_{NYC} – IERQ,DR) – ARR
Proposed Approach	when directing	ARR>CR or DR>CR	Q_{NYC} – IERQ – max(ARR,CR)

The *when to apply* condition was updated to apply the balancing adjustment when directing at any amount, not just those above conservation release (CR). Doing so provides unrestricted opportunity for the cumulative error to be resolved.

For the *when to accumulate* condition, the current approach accumulated error only when the directed release was greater than the conservation release. This approach had the goal of not accumulating error when NYC was releasing more than directed to maintain conservation flows, which are independent of the directed release requirement. However, this approach did not address the situation in which DR<CR but ARR>CR, for which error should also accumulate. The additional condition ARR>CR was therefore added to the algorithm. Also, the current approach accumulated only when directing. It was decided that accumulation of the error should also occur on days when ODRM did not direct (DR=0) but should have (ARR>0), so “when directing” was dropped from this condition, as indicated in Table 5.

In examining the *what to accumulate* term, the current approach is to accumulate min(Q_{NYC}-IERQ,DR)-ARR, but this is not appropriate when DR< Q_{NYC}-IERQ. An example is when a directed release is revised to be lower than originally computed due to an increase in the rainfall forecast. However, due to physical constraints of valve operations, timing of reduced releases, as well as Rapid Flow Change Mitigation protocol in the 2017 FFMP, a greater amount is released than needed from the reservoirs. In the scenario DR<ARR< Q_{NYC}-IERQ, the current approach would decrease the cumulative error (in the direction of "the reservoirs owe the river") by the negative amount DR-ARR when, actually, New York City released more than was needed. A more appropriate approach was determined to remove DR and incorporate CR into the equation and accumulate (Q_{NYC} – IERQ) – max(ARR,CR). Doing so ensures that the error is calculated using the actual release and is in the appropriate direction with respect to the actual required release and the conservation release. When CR>ARR, the error is calculated with respect to CR so that releases in excess of CR contribute to the cumulative error. The term Q_{NYC} – IERQ is always equal to or greater than CR, so negative errors do not occur when CR>ARR. When ARR>CR, the error is calculated with respect to ARR, i.e., the release that was actually needed to meet the Montague target flow. In this case, the error may be positive or negative. These scenarios what to accumulate are illustrated in Figure 7.

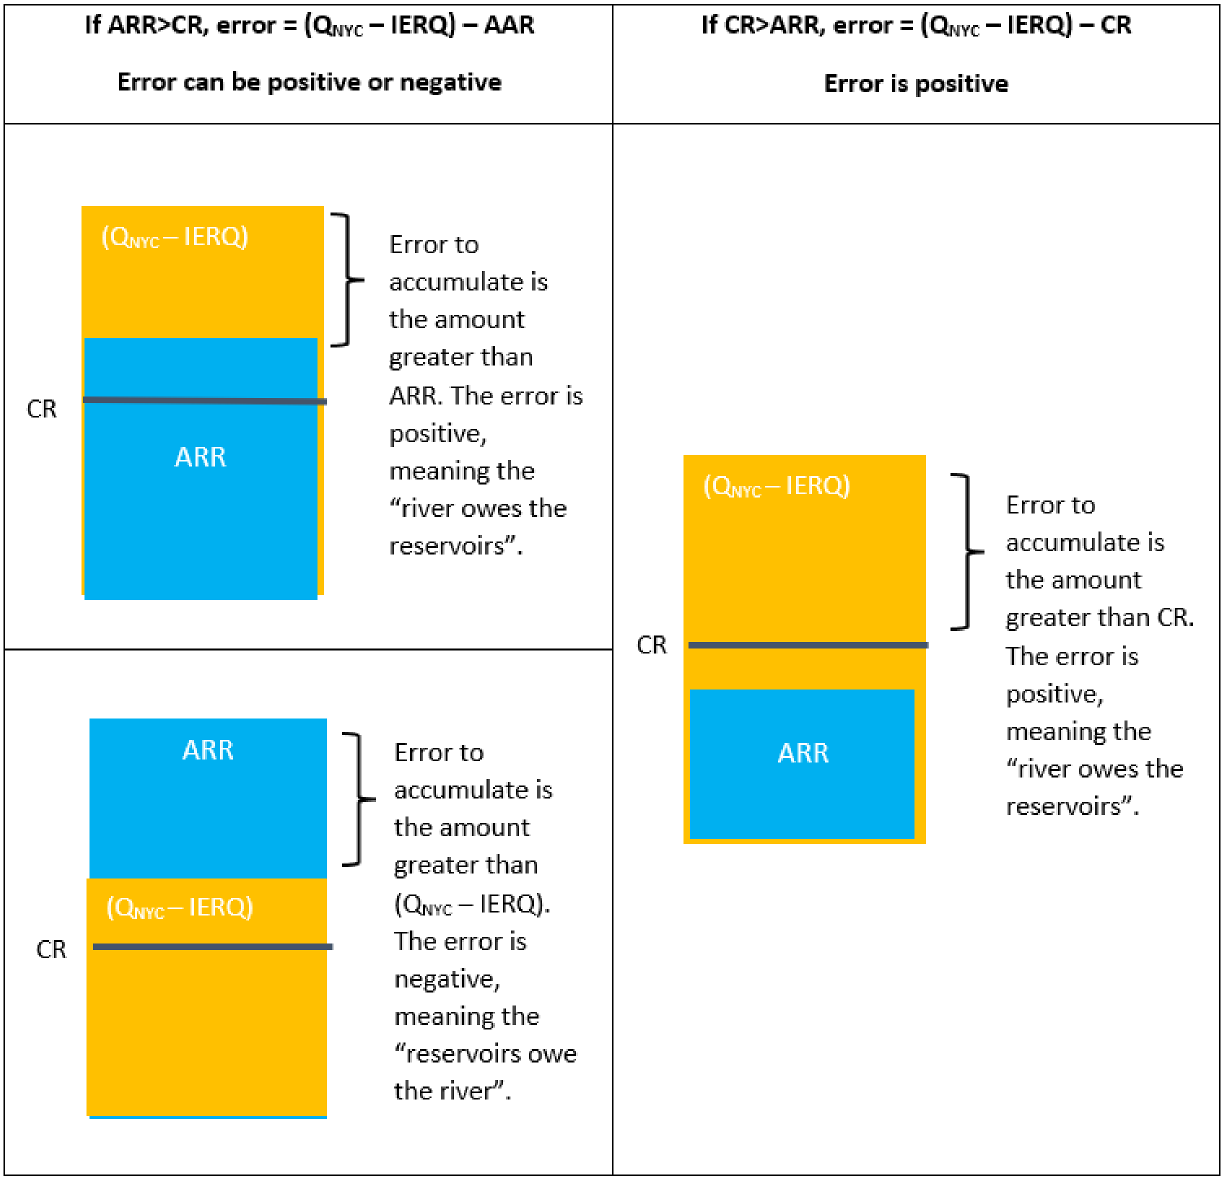

Figure 7. Illustration of the error to be accumulated for the balancing adjustment. The error is computed when either directed releases are greater than conservation releases, or when the actual required releases are greater than the conservation release. [ARR, Actual required release; CR, Conservation release; IERQ, Flow used for IERQ banks; Q_{NYC}, Total release from Pepacton, Cannonsville, and Neversink Reservoirs.]

Note that the BA equation takes the difference of two cumulative sums, which is equivalent to taking the cumulative sum of the differences. The difference being accumulated is the *what to accumulate* quantity in Table 5. In the case of the current approach (Table 5, top row), the cumulative sums are of the NYC release in response to the ODRM DR, calculated as min(Q_{NYC} – IERQ,DR), and of the actual release required (ARR). The equivalent sum of differences would therefore accumulate the term min(Q_{NYC} – IERQ,DR) – ARR each day, as shown in Table 5.

This cumulative error was calculated for the study period for the current and proposed approaches of Table 5. A third approach was also included for reference wherein no balancing adjustment was applied to the DR, but the resulting error was still tracked. The results for these three scenarios are shown in Figure 8.

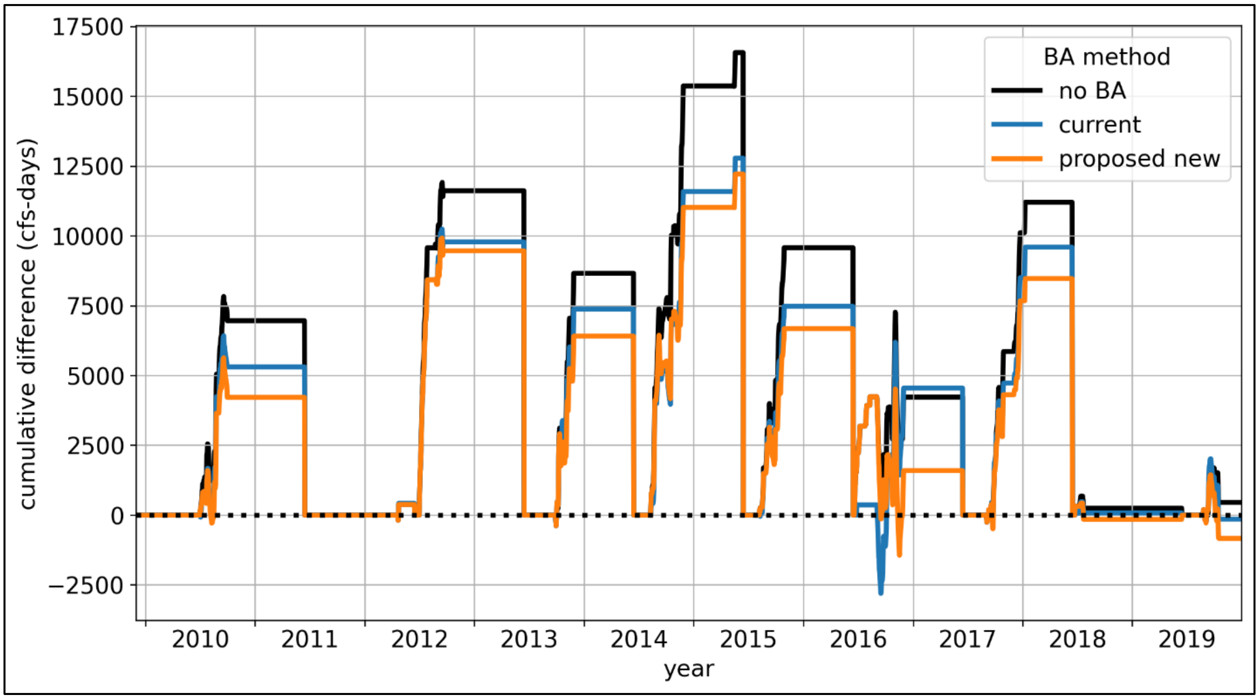

Figure 8. Cumulative error for different balancing adjustment methods.

Figure 8 shows that the current approach provides a better result than applying no BA in terms of maintaining a smaller cumulative error and that the proposed new approach is generally a little better than the current approach. The goal of the proposed new approach was not necessarily to improve this balance metric; the intent in developing this new approach was focused on closing loopholes and meeting the intended purpose of the balancing adjustment. Nonetheless, it is helpful to see that it does not degrade performance in terms of the cumulative error.

Note that changing the *what to accumulate* term changes the balance metric being used to evaluate the different approaches. Comparisons of such a metric among the different scenarios would not be meaningful. Therefore, in order to compare "apples to apples", the balance metric being plotted in Figure 8 is calculated separately from the cumulative error used to determine the BA (Table 6). The latter is calculated using the criteria given in Table 5 (compared to Table 6, left column) and the obvious choices for the "no BA" scenario. A second metric for the sole purposes of plotting and comparison, shown in Figure 8, is derived by always accumulating the same quantity, regardless of the scenario. That quantity accumulated is from the proposed new approach in Table 5, namely Q_{NYC} − IERQ − max(ARR,CR). This quantity was used because, as discussed previously, it was determined to provide the fairest, most logical representation of the cumulative error.

While the *what to accumulate* term was the same when calculating the cumulative sums used for plotting and comparison across all three scenarios, the *when to apply* term varied according to the scenario, as shown in Table 6. For the "No BA" scenario, the choice of which *when to apply* condition to use is somewhat arbitrary. Therefore, conditions for *when to apply* from each of the other two scenarios were evaluated and yielded similar results; only the one using the condition from the "Proposed Approach" scenario is shown in Figure 8. The right column of Table 6 shows how and when the balance metric used for plotting and comparison was accumulated for each of the three scenarios.

Having two separate balances calculated in these scenarios, one that accumulates the same quantity across all scenarios, and one that accumulates the quantity specific to that scenario, allows us to have a consistent metric with which to compare results while still applying the different approaches to calculate the BA required for each scenario.

*Table 6. Approaches used in the scenarios for calculating the two balance metrics—one for use in calculating the BA, and one for plotting and comparison (Figure 8). The rows in each entry are 1) what to accumulate, 2) when to accumulate, and 3) when to apply. *Using the "proposed" and "current" approaches for "when to accumulate" in the "No BA" scenario for plotting yielded similar results; only the former is shown in Figure 8.*

	Calculating BA	Plotting and Comparison
No BA	1) irrelevant 2) never 3) never	1) Q_{NYC} – IERQ – max(ARR,CR) 2) ARR>CR or DR>CR * 3) not applied
Current Approach	1) min(Q_{NYC}-IERQ,DR) – ARR 2) when directing and DR>CR 3) when directing and DR>CR	1) Q_{NYC} – IERQ – max(ARR,CR) 2) when directing and DR>CR 3) not applied
Proposed (revised) Approach	1) Q_{NYC} – IERQ – max(ARR,CR) 2) ARR>CR or DR>CR 3) when directing	1) Q_{NYC} – IERQ – max(ARR,CR) 2) ARR>CR or DR>CR 3) not applied

<u>Reset</u> – Model results showed that resetting the cumulative error on various dates instead of the current standard of June 15^{th} had no effect on BA.

A suggestion was made to reset the cumulative error when reservoirs spill, or flows were above a given flow amount (the arbitrary value suggested was 5,000 ft^3/s). The suggestion was based on the idea that the reservoirs and river were considered to be replenished naturally under such conditions, and therefore no compensation due to directed release error would be needed going forward from that point. However, the idea was countered on the basis that reservoir spills rarely occur during the summer and late fall when directed releases are needed, and that the 5,000 ft^3/s value was arbitrary. The option was not considered to be judicious did nor did it serve to simplify the process.

Feedback from the workgroup indicated that January 1^{st} was not desirable as it was arbitrary related to flow management. The suggestion of shifting the date from June 15^{th} to June 1^{st} appeared reasonable, as it would align with reset dates for banks used for the Interim Excess Release Quantity.

<u>Distribution</u> – Analysis of the model results indicated that the historical approach of distributing the BA by dividing the cumulative error by 10 had limited benefit. It effectively causes the cumulative

error to be unresolvable because the resulting balancing adjustments were too small, and it is also somewhat redundant given the BA cap, which was set at 50 ft^3/s. Feedback from the workgroup concurred with the removal of this component of the balancing adjustment calculation.

Caps – The parameter that has the greatest impact on resolving the cumulative error is the cap on the absolute value of the balancing adjustment. To see how varying this parameter affects the cumulative error, multiple pyDR runs were performed, each with the cap parameter taking on a different value (Figure 9). For these and all subsequent results, the "proposed approach" delineated in Table 5 (second row) is used. The "no BA" results from Figure 8 are retained for reference.

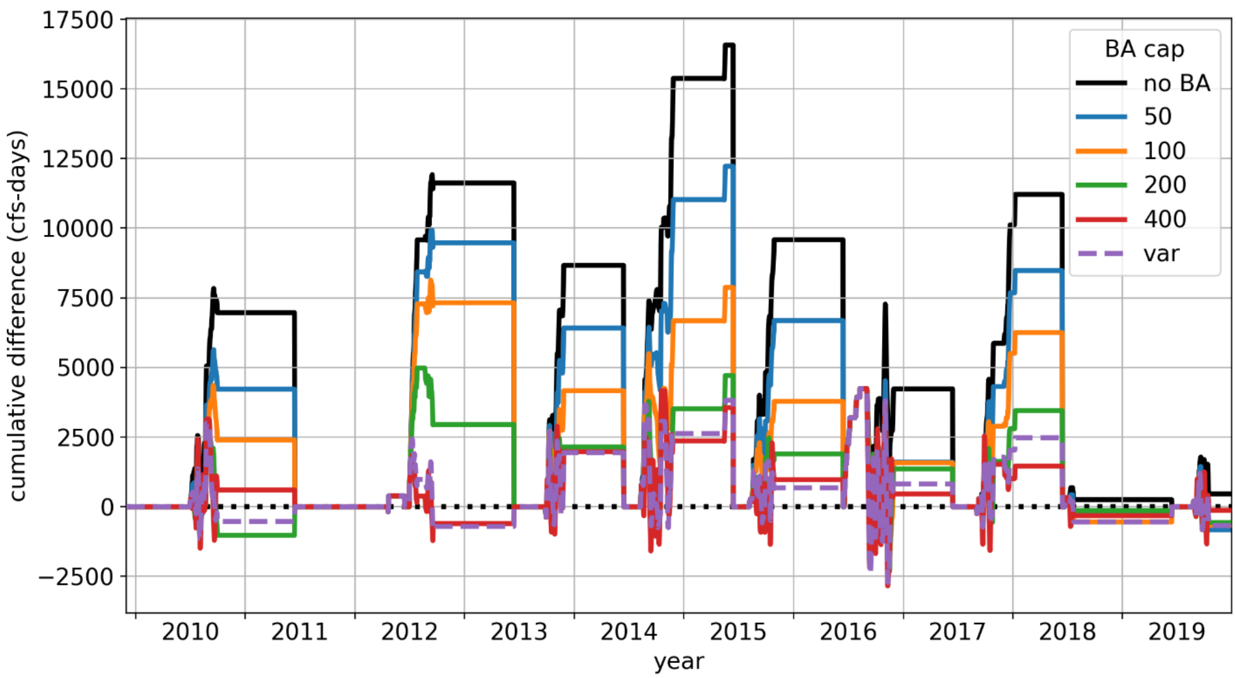

Figure 9. Cumulative error for different balancing adjustment caps.

In addition to the range of numerical values for the cap (50, 100, 200, and 400 ft^3/s), an additional approach labelled "var" was evaluated. This approach is discussed further below.

Clearly, larger caps result in lower balances, as would be expected. A tradeoff of achieving the lower balances is that Montague flows tend to deviate more from the flow objective (which is generally 1,750 ft^3/s). Figure 10 shows simulated Montague flow for the same scenarios of Figure 9 for two example years, 2014 and 2016.

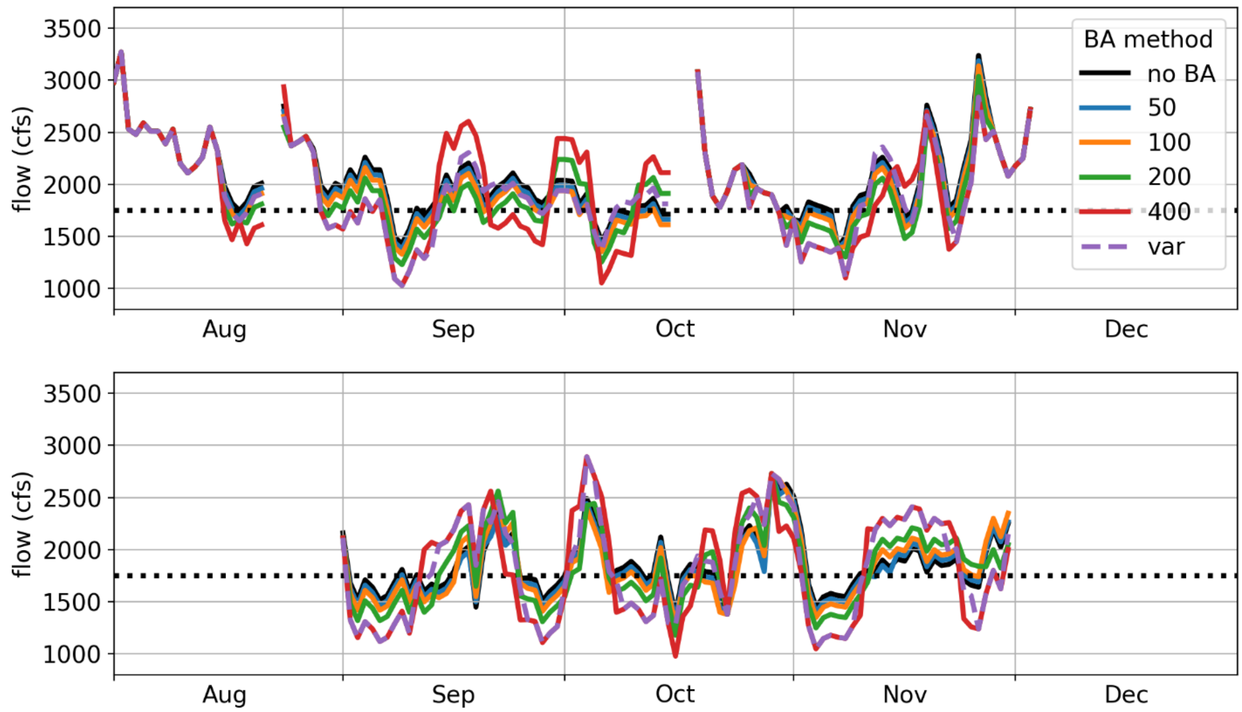

Figure 10. Comparison of Montague flows using different BA caps in example years 2014 (top) and 2016 (bottom). The target flow is shown as a dotted line.

Like the two years shown in Figure 9, all years (not shown) portray higher caps resulting in larger deviations from the target flow, an undesirable outcome. This is demonstrated more explicitly using the flow probability density functions for the scenarios (calculated using the kernel density estimation method, Rosenblatt 1956 and Parzen 1962), shown in Figure 11. Generally, higher caps result in a wider distribution, with fewer days having flows near the flow target and more days with flows farther removed from the target.

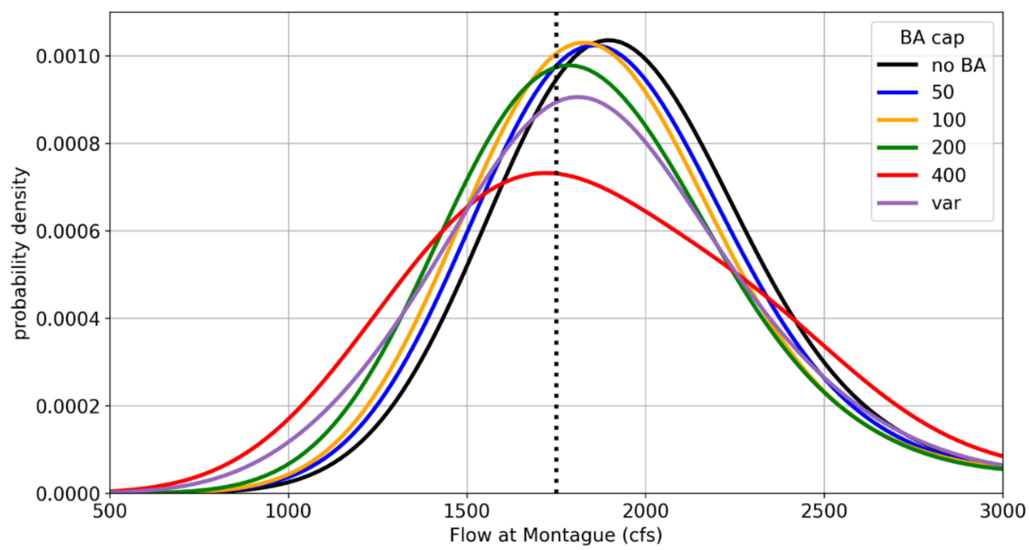

Figure 11. Flow distributions for different balancing adjustment caps. The dotted line is the dominant flow target, 1,750 ft^3/s

Taken together, Figure 9 and Figure 11 illustrate an apparent tradeoff between resolving the cumulative error and minimizing deviations from the target flow. In an attempt to better balance these goals, the “var” flow cap was proposed and evaluated. In this approach, a variable cap is used wherein the cap is 400 ft^3/s when the cumulative error is >=1000 ft^3/s-days and 100 ft^3/s otherwise. This implementation of the larger cap only when the cumulative error is large reduces the overall risk of large flow deviations from the target. In terms of the cumulative error, the variable-cap approach performs very similarly to the 400 ft^3/s flow cap (Figure 9), i.e., achieving the same benefit of a reduced cumulative error. In terms of flow deviations from the target, the variable-cap method is a clear improvement over the 400 ft^3/s cap, being more similar to the 200 ft^3/s cap with reduced occurrence of large deviations from the flow target. This variable approach is therefore a promising alternative and could in theory be adapted to use any maximum cap, reducing the cumulative error while minimizing the increase in flow deviations.

The majority of the feedback during the July 2022 workshop indicated that the balancing adjustment was needed and should not be removed completely. The majority also agreed that a cap of 100 ft^3/s was an acceptable change to make the balancing adjustment more efficient, though some support was given to have it remain at 50 ft^3/s. Caps of 200 ft^3/s and greater were viewed as unacceptable, including the variable cap which was assumed to increase complexity. Being overly complex was one of the original criticisms of the procedure.

Final Recommendations

As a result of the review of the data, model output, and feedback from workshops, the following modifications to the Balancing Adjustment Procedure were recommended. The changes were implemented beginning June 15, 2023, with the exception of Item 1c, which will be updated in 2024.

1. Simplify and increase effectiveness:
    a. Increase the cap from 50 to 100 ft^3/s.
        i. Increasing the cap will help to resolve the cumulative error more efficiently.
    b. Remove the distribution (divide by 10) and rely only on a maximum value (cap) for the Balancing Adjustment value.
        i. Dividing the cumulative error by 10 was found to overly distribute the error to the point that it was unresolvable. It is also duplicative of using a cap.
    c. Change date from June 15 to June 1.
        i. Align with IERQ banks and diversion calculations reset dates.
2. Correct loopholes in algorithm:
    a. When to apply: Apply the Balancing Adjustment not only when directed releases are greater than conservation releases, but for any directed release value.
    b. When to accumulate: Continue to accumulate error not only when directed releases are greater than conservation releases, but also when the actual required release is greater than the conservation release. Doing so will include errors when the directed release was less than conservation but should have been greater than the conservation release.
    c. What to accumulate: Accumulate error as the difference between the NYC reservoir release in response to ODRM and the maximum of the actual required release and conservation release.
        i. All corrections to the algorithm will ensure that the occurrence and magnitude of forecast errors will be appropriately accounted for in relation to the conservation release.
3. Continual tracking and improvement:
    a. Review past ODRM reports to build a history of the Balancing Adjustment. For example, it was not always capped at 50 ft^3/s.
    b. Explore effects of the balancing adjustment during drought (eg, 1960s, 2001-02). The ODRM 5-year plan for fiscal years 2024-2028 includes an ODRM data retrospective, which can include a focus on the Balancing Adjustment Procedure.

References

Parzen, E. (1962). "On Estimation of a Probability Density Function and Mode". The Annals of Mathematical Statistics. 33 (3): 1065–1076. doi:10.1214/aoms/1177704472

Rosenblatt, M. (1956). "Remarks on Some Nonparametric Estimates of a Density Function". The Annals of Mathematical Statistics. 27 (3): 832–837. doi:10.1214/aoms/1177728190

Appendix 1. 2023 Balancing Adjustment Study Statement of Work

Introduction

An empirical model, incorporating a recession and routing strategy, is used for forecasting the average daily flow at Montague. ODRM segregates the previous day's flow at the gaging station to calculate flow in the river from uncontrolled input (ie. not from NYC reservoirs or powerplants). ODRM uses forecasts of powerplant inflows and runoff from rainfall to determine whether releases are required from New York City reservoirs to maintain the flow objective at Montague. These procedures have been refined over nearly six decades of use and generally give good results during periods of relatively stable flow.

However, as stated in the ODRM 1958 annual report, *"It is impossible, because of the many factors involved, to precisely predict in advance the resultant flow at Montague. However, over a period of time any cumulative error in the estimating procedure should be eliminated as far as possible. To accomplish this, a balancing adjustment is applied to the indicated deficiency or excess in computing the directed release. This adjustment is based on the* ***amount by which the directed release is greater or less than the release actually required to maintain the prescribed rate of flow at Montague****…."* This was the first instance that the balancing adjustment was described in an annual report and ODRM has utilized different procedures to calculate the amount since that time.

Currently, the Balancing Adjustment is applied only when the directed release is greater than the conservation release. Conservation releases from the New York City reservoirs are specified in the 2017 Flexible Flow Management Program. No Balancing Adjustment calculation occurs if the directed release is less than the conservation release sum from Cannonsville, Pepacton, and Neversink Reservoirs. This approach prevents accruing a Balancing Adjustment during times where the conservation releases to the Delaware River are greater than the directed releases required. When the directed release is larger than the conservation release, the Balancing Adjustment procedure is applied.

When directed releases are determined to have been excessive (greater than required), then the Balancing Adjustment reduces the directed release requirement. Conversely, if prior directed releases have been insufficient, then the Balancing Adjustment increases the directed release requirement. The Balancing Adjustment is computed as 10 percent of the difference between the cumulative directed release and the cumulative directed release required for exact forecasting, and is applied until the difference reaches zero. Using a percentage of the cumulative difference helps to distribute correction over time. The maximum daily balancing adjustment is also limited to +/- 50 ft^3/s to further preclude unacceptably large variations in the adjusted flow objective.

Objective

Section IV.4 of the Flexible Flow Management Program 2017 (FFMP 2017), requires the Decree Parties to study, evaluate, and consider the River Master's balancing adjustment procedure. The scope of work below was developed in order to accomplish this study.

Approach and methods

Directed releases data from 2010 to present will be compiled and summarized. A group of experts will then review the data and determine possible alternatives that could improve the process. The possible alternatives will be tested, reviewed, and presented to the Decree Party Work Group. A decision to implement a new or modified Balancing Adjustment will be made by ODRM in coordination with the Decree Parties, and a report documenting the study and outcome will be produced. Details of this approach are listed below:

Tasks

1) Data Collection and Model Development

1. The ODRM directed releases information has been compiled from January 1, 2010, to December 31, 2019. This will be utilized as the input information in the scenarios. However, this data included the current process of correcting for the balancing adjustment. The compiled file will be modified to remove the influence of the current process.
2. The current script that was developed for the design modernization process will be utilized again for this study. Some updates were requested at the January 7, 2020, workshop meeting and these will be incorporated.
3. Current process analysis. The current process will be analyzed to see how the balancing adjustment performed over the period of data. Some pre-existing criticisms that can be evaluated include that the current process is "overly complex", "isn't effective" and "carries a balance for long periods of time". Performance in dry vs wet years will also be analyzed.

2) Alternatives and Metrics Development

1) A workshop will be held to brainstorm alternative balancing adjustment methodologies and metrics for evaluation. Attendees will include USGS experts that have participated in the design modernization workshops and the Decree Party Workgroup (DPWG) (tentatively schedule around end of June timeframe)
    a) Alternatives for the methods regarding implementation of the balancing adjustment will need to be articulated and then distilled to a list for consideration and testing.
    b) Metrics that will test the performance of the balancing adjustment methodologies will be brainstormed and the selected for assessment.
        i) Current ideas include, but are not limited to:
            (1) How often through the year is the balancing adjustment at 0?
            (2) How frequently is the balancing adjustment capped at +50 or -50 ft^3/s?
            (3) What is the magnitude of cumulative excess or deficit that the current balancing adjustment approach results in each year?
2) The set of alternatives and metrics will be shared with the public for input and feedback (likely through DRBC's Regulated Flow Advisory Group (RFAC)).

3) Scenario Performance Testing

1. The alternatives selected during the workshop will be tested with the script and selected metrics.
2. Results will be shared with the DPWG for feedback and potential further evaluation of no more than three alternatives.
3. Final results will be shared with the DPWG and the public (likely through RFAC).

4) Reporting and Decision

1. A report will be produced that explains the methods used to evaluate the alternatives and the results.
2. A selection of the new methodology for the balancing adjustment will be made by the Office of the Delaware River Master in coordination with the Decree Parties.

Coordination and Communication

A Design Modernization Study (DMS) will be occurring concurrently with the Balancing Adjustment Study. One goal of the Design Modernization is to improve forecasted flow at Montague. Inputs from the DMS may necessitate changes to the Balancing Adjustment Study scope after work has begun. The balancing adjustment team will coordinate with USGS experts that have been engaged in the ODRM design modernization process. At least one workshop will occur with this group to develop different methods to incorporate a balancing adjustment and metrics to evaluate these methods. These USGS experts are already engaged in the directed release process and know how best it can be improved, including the balancing adjustment.

In addition, the DPWG will be invited to the same workshop to provide input into the different methods and metrics. There will be regular updates to the DPWG and opportunities for input and feedback. Meetings with RFAC will occur to solicit public input and present study results.

Team

Kendra Russell – ODRM

Noah Knowles – USGS

Liz Hittle – USGS

Vin Difrenna – ODRM

Amy McHugh – ODRM

Schedule

Activity	Completion Date	Adjusted 12/13/21
Input data modification	June 1, 2021	
Current process analysis	June 1, 2021	
Prep for workshop meeting	June 15, 2021	
Alternatives and Metrics Workshop	June 30, 2021	
Public Input on Alternatives and Metrics Selection	RFAC Meeting	
Script adjustments	August 1, 2021	October 31, 2021
Scenario Testing	August 31, 2021	December 31, 2021
Results Workshop	October 31, 2021	January 31, 2022
Draft Report	February 28, 2022	June 30, 2022
Decision/Changes implemented	May 30, 2022	September 30, 2022

Balancing Adjustment Modifications Discussed at DPWG Meeting in March, 2019

- Remove cap of +/- 50 ft^3/s
- Implement a threshold under which the entire balancing adjustment is resolved daily (rather than dividing out over 10 days)
- Have the balancing adjustment released at +/-50 ft^3/s until it is used up (ex 200 ft^3/s in 50 ft^3/s increments for 4 days)
- Have a larger cap than +/- 50 ft^3/s
- Use balancing adjustment as a bank
- Does the balancing adjustment need to reset?
- If it resets, have the balancing adjustment reset aligned with another “year”
- No balancing adjustment

Appendix 2. 2023 Balancing Adjustment Calculation Flowchart

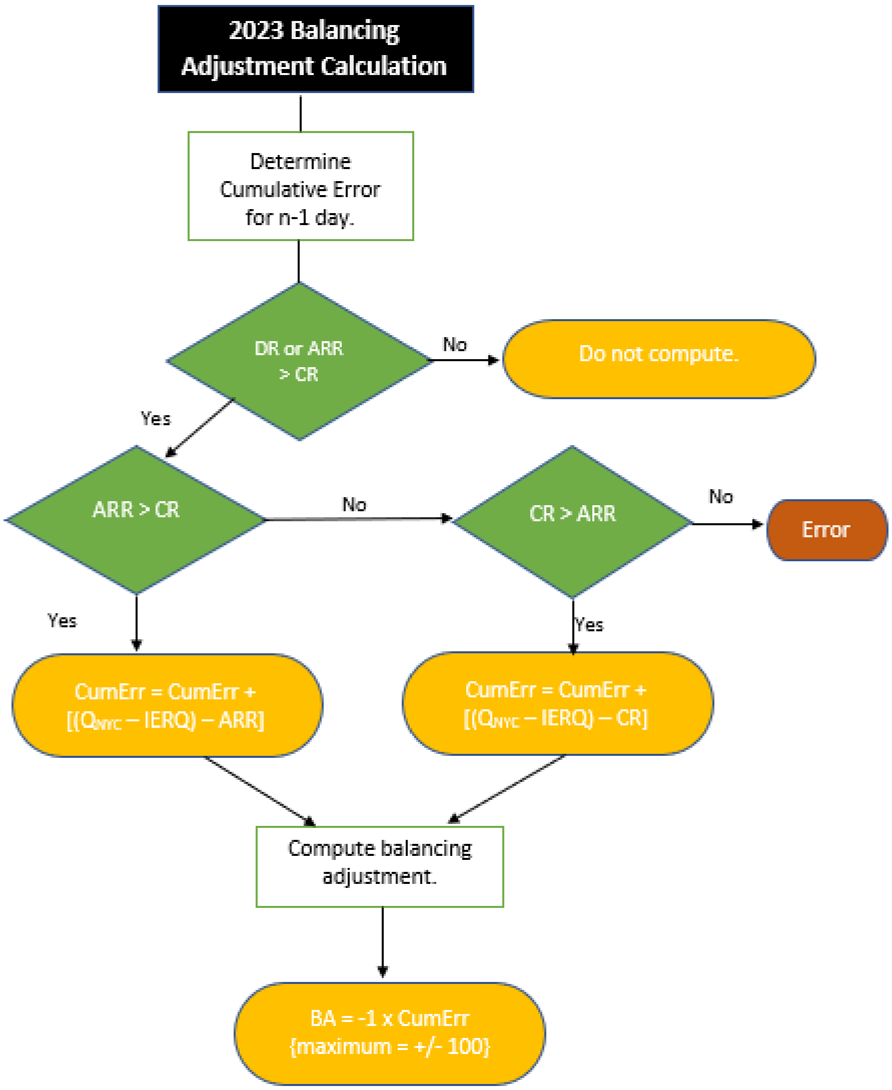

Terms and definitions:

ARR	Actual required release: Flow Objective – (Previous day's daily mean discharge at Montague – Q_{NYC} – IERQ), >0 only
BA	Balancing adjustment
CR	Conservation release
CumErr	Cumulative error with respect to conservation release; CumErr = 0 beginning each June 1.
DR	Directed release
IERQ	Flow used for IERQ banks
Q_{NYC}	Total release from Pepacton, Cannonsville, and Neversink Reservoirs