Bay Oxygen Research Group (BORG) Meeting

December 15th, 2025
12:00 PM – 1:30 PM
Visit the meeting webpage for meeting materials and additional information.

Purpose: This is the December 2025 Bay Oxygen Research Group Meeting. During this meeting, Rebecca Murphy (UMCES) presented how data is used within the 4-d interpolator by drawing from the presentations given by her and Jon Harcum (TetraTech) at the Coastal Estuarine Research Federation Conference. This presentation sparked many questions, conversations and future action items within the group including the use of and methods for community science data, spatial correlation, and smoothing curves.

Minutes

I. Welcome, Introductions & Announcements

Lead: ***Breck Sullivan*** *(U.S. Geological Survey, USGS)*

Follow-up from November BORG Meeting

Breck dove into the decisions and actions from the most recent BORG meeting to provide updates.

1. Proposed thinning:

While developing the 4-d interpolator, the development team wanted to ensure that users could run the tool on their laptops rather than needing a larger set-up. In testing this, five of the segments were unable to run reliably on a laptop. The Development Team proposed thinning these to have larger resolutions and the group agreed on this. The group also brought up the idea of thinning the other 50m resolution segments to 100m. It was decided that the Development Team would test the other 50m resolution segments to ensure the quality of the results would remain the same. The Development Team will run this testing soon and bring this topic back to the group.

2. Subsampling and weighting:

At the November BORG Meeting, Rebecca brought up an issue they were facing when using all of the high frequency data. The high frequency data was disproportionately impacting the results at nearby stations, which is known as “imbalance of classes.” To overcome this, the group proposed sub-sampling the data to an hourly frequency, rather than 15 minutes. When the sub-sampling is conducted, there is little to no change in the distribution of the data, meaning nothing is lost that would change criteria assessment conclusions. The members had some suggestions for whether all data should be incorporated and at what frequency. The Development Team is going to continue researching this topic. Members also suggested looking into weighting the data which

the team is now testing. This conversation will be brought back to a future discussion to make a final decision on the matter.

Breck wants members to know that both of these topics are top of mind for the Development Team and once they have collected enough information and preformed the tests, they will be sure to bring the results and conversations to a later meeting.

II. How Data is Used in the 4-D Interpolator

Lead: ***Rebecca Murphy*** *(University of Maryland Center for Environmental Science, UMCES) and* ***Jon Harcum*** *(TetraTech)*

Rebecca and Jon pulled together slides from their presentations at the Coastal Estuarine Research Federation Conference and presented them at this meeting to help members understand how the tool uses data. Currently, one of the goals and difficulties with the interpolator is using all of the available monitoring data. The interpolator aims to use multiple types of data to obtain Bay-wide estimates of dissolved oxygen concentration in horizontal, vertical, seasonal, and within-day. They use a generalized additive model (GAM) to calculate the mean mid-day DO as an interpolation everywhere and every day. To balance the fixed station data that has a lot of variability in when it's observed, they put it into a daily estimate and then add correlation terms and daily cycles that are fit to the observed data. Rebecca presents all the types of data that is being used in the interpolator and for what step it is being used.

Rebecca explained how the mean mid-day space-and-time interpolation is created and explained which data is used for this step, which is all of the data in each segment sub-sampled to an hour. Being sub-sampled to an hour is not a firm decision but is working well, as discussed in the last meeting. Each segment uses data from boundary segments, which border the original segment.

Next, Rebecca dives into the correlation data with an example on horizontal correlation. Using DataFlow data, the group created variograms. These variograms found that at a close distance there is high spatial autocorrelation, but as you get further away, it decreases. Elgin addressed that the variograms level off at a certain distance which could point to the fact that data 3km apart is not good at estimating the 1km correlation coefficient.

Next, Rebecca presented a map of the mean 1000m correlation throughout the Bay, which showed that most means fall within 0.4 and 0.7 and that horizontal correlation is higher in deeper waters. A smooth spatial function fit to these correlations will use location and the correlation computed from the variograms to estimate horizontal correlation everywhere. Lastly, Rebecca shared how these steps would create the hourly realizations for the tidal waters of the Bay.

Actions:

1. Do research on whether there is continuous dissolved oxygen in Tier 2 Chesapeake Monitoring Cooperative data to support correlation computations and discuss with the group how Tier 2 data can be used to support the tool even if it is not used for criteria assessment.
2. Discuss the schedule and feasibility of updating the correlation computations using new data as available going forward in operational mode.
3. Investigate using the observed data along with the simulations when utilizing the cumulative frequency diagram (CFD) curves. Future work could also be to enhance the 4-d model to nudge it to the observed data, but first ensure smoothing curves are correct.
4. Moving forward with comparisons between the 3-d and 4-d interpolator once the 4-d interpolator is complete.

Discussion Notes (during presentation):

Comment (in chat): *Jim Hagy:* When you talk about vertical correlation, I think you also mean not just similarity between two adjacent depths, but if dissolved oxygen (DO) is lower than expected at one depth, than it's also lower than expected at nearby depths.

Q: *Matt Stover:* This is very helpful. I appreciate how you're explaining this. You mentioned that the DataFlow is used for interpolating horizontally. Why wouldn't you use fixed stations for that? Especially if fixed stations or ConMons were reasonably close to each other. I know for DataFlow you're getting data every five seconds which is spatially dense, but there is a time aspect. If you're running up one side of a tributary, it might be an hour or two until you run down the other side.

- ***A:*** *Rebecca Murphy:* We address the time part in how we split the data when we build variograms and get the spatial correlation. Like you said, we don't want data that's an hour apart to be grouped together. It's not. You'll see a map. We didn't use the other data sets for this because we need to parameterize the spatial correlation at 1km. There are probably some data points that are 1km apart, it's not that common. We didn't pull that out. It probably wouldn't nudge things too much.
- ***Response (in chat):*** *Mark Trice:* Severn River. Re: dataflow DO differences from time going up one side then back down the other.
- ***Response:*** *Jon Harcum:* You hit that on the head. There isn't a large data component, besides those identified, that would help us with the spatial correlation. If you know of a data set, please point me to it and I'll be happy to take a look and add it to our analysis.
- ***Response:*** *Matt Stover:* I was imagining that if you have a hypothetical tributary and there are four or five fixed stations collecting DO at multiple depths. I think you'd be capturing the layer cake of the tributary. They might be further than a kilometer apart, but you could learn something from the spatial differences in the DO at different depths.
- ***Response:*** *Rebecca Murphy:* That data would be taken into account when informing the mean mid-day space-time interpolation. We have the spatial relationship as you go up the Patuxent River at a certain depth, for example.
- ***Response:*** *Elgin Perry:* Your thinking is right on target. That data would be very useful. If we had a sequence of vertical arrays that are 1km apart and went surface to bottom, we could get depth correlation and horizontal correlation below the surface. If that data was around, it would be very useful.
- ***Response:*** *Matt Stover:* We had previously talked about where to put arrays. For example, if we put five arrays equidistant along the Severn River, which is 10-15 miles. Would that be useful? You'd be getting multiple depths at the same time, where DataFlow varies in time.
- ***Response:*** *Elgin Perry:* You're right on target. We get horizontal correlation through a variogram analysis, which parameterizes correlation as a function of distance. If they were 2km apart and we wanted the correlation 1km, we'd just take the square root of the correlation. We've already had to do that in the depth dimension because the arrays are traditionally 2km apart. That could be a way to fill in the horizontal correlation below the surface. Another thing is a scan fish device. You'd tow it behind your boat, but it also has the ability to dive so you can hold it at a certain measure of depth and measure a series of depth observations. We could get horizontal correlation from that as well. There are a couple ways to get at that parameter, but we don't have the data to do it yet. The way you are thinking is just the way to get that.
- ***Q:*** *Matt Stover:* Is it that the beyond a kilometer the correlations aren't good? Are they not helpful at longer distances because of changes in depth and bathymetry?
- ***A:*** *Elgin Perry:* With fixed station data, we are talking 10s of kilometers. That is not that helpful. Beyond a kilometer could be ok but beyond 5km would not be helpful. I'd rather have the data and let it tell me what the cutoff is but at this point we don't have the data.

- ***Comment:*** *Matt Stover:* We've been thinking about Fishing Bay. We've blanketed it with stations. I don't know if that's a repeatable sampling model. When filling in the MDE shellfish station monitoring, does that give us better horizontal correlation?
- ***Comment (in chat):*** *Jim Hagy:* Having equally spaced stations can leave the left side of a variogram looking a little thin. You can model as Elgin suggests, but you really need either randomly spaced samples, or continuous underway data.
- ***Comment:*** *Mark Trice:* Twenty years ago, I looked at correcting DataFlow using adjoining ConMons in a segment. The Severn was particularly noticeable where the data changed a lot between going up the River in the morning and down later in the day. The ConMon signal was used as a correction factor. It might make sense if a process was developed to do that formally. You'd be inputting time corrected DataFlow data into the model.
- ***Response:*** *Elgin Perry:* We have a process in place. We've opted to look at the boat speed and we subset to data where the boat is moving quickly. Then we look at a window of distance along the DataFlow string. We do our variogram analysis on data points that are within a distance along the string. Since the boat is moving fast we feel like they're time matched. I did some work on using local ConMons to adjust the time signal of the DataFlow, but I found it wasn't successful sometimes. I feel more comfortable with the windowing approach we are using. We can discuss that in the future too. We have thought about that and that's the solution we came up with.
- ***Response:*** *Rebecca Murphy:* On slide 4, we show a map of the DataFlow. The different colors are the splitting of the DataFlow into individual units so the information used to build variogram and the correlation structure would only be from these segments. This way we wouldn't be having these concerns. It would be a shorter stretch of the DataFlow.

Q (in chat): *Melinda Culter:* Is there a reason why 1km / DataFlow data was selected then, if mathematically we can use the variogram to include stations from farther distances?

Q: *Tish Robertson:* I am wondering if there is value in looking at the Tier 2 citizen monitoring data when calculating the relationships. It makes sense that we'd restrict it to Tier 3 for the actual assessment, but if we're just looking at spatial or temporal patterns, it seems like Tier 2 would be useful.

- ***A:*** *Rebecca Murphy:* We didn't use it for any of the correlation parts because the citizen monitoring is more like going out once a week at one depth. It doesn't fill the needs for any of these correlations. Although, I think some of the programs have a couple of continuous stations. If we want to use Tier 2 data for the mean mid-day, we can. The structure is there to add it to the interpolation.
- ***Comment:*** *Tish Robertson:* I was throwing that out as an idea for the group to consider. We could have it as a potential input to think about. Tier 2 isn't good enough for assessment, but that doesn't mean it's not useful, so it might be worth it for the group to consider how we could use it in development.
- ***Response:*** *Rebecca Murphy:* I agree. As we run an interpolation and see how well it's performing, that would be great data to use for comparison.

Q: *Leah Ettema:* The mid-day space-time interpolation is done at every centroid?

- ***A:*** *Rebecca Murphy:* It's done in a grid throughout an assessment unit. An assessment unit is a segment of the tidal waters and there would be a grid that is maybe 1km x 1km x 1m. We'd get a result in every one of those grids in a segment every day.
- ***Q:*** *Leah Ettema:* Is that what the current interpolator does too?
- ***A:*** *Rebecca Murphy:* Yes.

- ***Q:*** *Leah Ettema:* I am not familiar with the horizontal correlation in the 3-d interpolator. How is this different than the 3-d interpolator?
- ***A:*** *Rebecca Murphy:* The 3-d interpolator does not do temporal interpolation. The mid-day space-time interpolation is not the same between the 3-d and 4-d interpolator. The 3-d interpolator takes all of the data in the first half of June and interpolates it in space, which is the early June results. It does the same thing with late June.
- ***Q:*** *Leah Ettema:* It uses kriging?
- ***A:*** *Rebecca Murphy:* No, it uses inverse distance weighting (IDW). The whole structure is different. The correlation part is new. The current interpolation assumes the inverse distance weighting relationship between oxygen everywhere. This new approach lets the data tell us what the relationship is, which could be different in different areas. It won't be the inverse distance weighting relationship.
- ***Q:*** *Leah Ettema:* So, this data informs each spatial direction of that relationship. You don't have all of the data to do the horizontal correlation, like you were saying to Matt's question.
- ***A:*** *Rebecca Murphy:* That is one challenge. Another challenge is that the data to parameterize this is spotty in space and time. It's great but it's impossible to collect high frequency data everywhere, the programs cycle where these are over the years. Right now, in the Bay, there's ConMon in some of the tributaries, but not all of them. Over the last 20 years, there has been ConMon in every tributary. We wanted to use all of that data to parameterize the correlations to get high frequency estimates everywhere in the current year, using a database of historical data.
- ***Q:*** *Leah Ettema:* Is the horizontal correlation going to be assumed to be what the surface data is?
- ***A:*** *Rebecca Murphy:* Yes, it is right now.
- ***Response (in chat):*** *Tish Robertson:* IDW = 7th grade arithmetic, GAMs = Graduate school stats
- ***Response (in chat):*** *Peter Tango:* 3D interpolates in space. We accumulate the results across time (each cruise over a 3 year period) when we apply the CFD attainment analysis to inform a result.
- ***Comment (in chat):*** *Jim Hagy:* Matching the correlation structure = Priceless
- ***Response:*** *Rebecca Murphy:* Jim, thank you for clarifying. For vertical correlation, if the DO is lower than expected at one depth, then it will likely be lower than expected at the next depth too.

Q: *Melinda Cutler:* This is really helpful, especially this flow chart (slide 7). Is the development of the correlation going to occur iteratively every time there is new data? Or is this something we do once and then that's the model?

- ***A:*** *Rebecca Murphy:* You ask great questions. That is something we have been discussing. Right now, we've parameterized this with a database of existing data. We were just talking about this last week. We need to work with you all and figure out what the needs are.
- ***Response:*** *Elgin Perry:* Yes, we're wondering too. One difficulty is that the ConMon data is done on a rotational basis. We can look back at history and have a dataset that covers the entire Bay. If you add a year of data, it will only give you data for a few ConMons in a few tributaries. I don't know if the update period will be annually or less frequently. That's something we need to talk about and come up with a schedule. Yes, the parameters should be updated as we go along.
- ***Response:*** *Rebecca Murphy:* The mean mid-day interpolation would definitely be updated with all of the data in that year to create that year's interpolation. With the correlation, it might not move things much if we update yearly, but we understand that it would be of interest and would like to. Over the next few months, we'll be talking about process.

- ***Response (in chat):*** *Jim Hagy:* Intuitively, I think the correlation structure should be a fundamental property of the system, not something that changes quickly, even if the DO status shifts.
- ***Q:*** *Melinda Culter:* Would we ever remove historical data? It would just be adding additional data.
- ***A:*** *Rebecca Murphy:* I think that would be worth a discussion too. In Virginia, I heard Dave Parrish say that their DataFlow cycle is starting over again. If we had data from 2005 for a tributary and we got new data for 2025, we might want to remove the data from 20 years ago and only use the current data. We can also test this and see how different it is. That would be really informative to see if spatial correlation is changing over time. We're not sure about that but would love your thoughts.

Discussion Notes (after presentation):

Q (from chat): *Becky Monahan:* If there is actual data, are we going to use the values of the data, or of the interpolation? I understand interpolating for areas where there isn't data.

- ***A:*** *Rebecca Murphy:* Right now, we would use the interpolation results. There are methods to nudge model results to the data. We haven't implemented that, but I can understand this question. We've been working hard to get this as good as we can. We are certainly going to be comparing the data in every spot where it exists, and we've been doing that. Right now, it would be using the interpolation results. This could be a larger conversation for the future.
- ***Response:*** *Becky Monahan:* I know you said the last slide was a draft and that it had gotten better since then. There were some values that were much lower than the interpolator projected. If we are collecting high-frequency data, I'd assume that we'd use that available data, rather than the interpolator. I would rather use the data when it's available but understand using the interpolator in those areas that we don't.
- ***Q (from chat):*** *Jim Hagy:* What is the ratio of interpolated locations and times vs. observed locations and times? That might illustrate how much substituting actual data for interpolated values would change an assessment. I'm thinking that the number of observed measurements is TINY
- ***A:*** *Rebecca Murphy:* Yes, I understand and I want to be true to the data. Jim had a helpful comment here. Generally, it's pretty accurate. I am not sure how the handful of data points in the amount of hours in a year and locations would adjust the results.
- ***Q:*** *Leah Ettema:* How many grids do you have in an assessment unit?
- ***A:*** *Rebecca Murphy:* I don't know off the top of my head. I'd have to open a spreadsheet.
- ***Comment:*** *Leah Ettema:* So, Becky is saying to use the data if it is available.
- ***Response:*** *Becky Monahan:* Yes. In the grids where it's not available, we can understand how you'd be interpolating between nearest grids that have data. We wouldn't want to interpolate if we have existing data.
- ***Q (from chat):*** *Peter Tango:* That would be consistent with the existing interpolation of data in the 3D for assessment, correct?
- ***A (from chat):*** *Tish Robertson:* Yes.
- ***A:*** *Breck Sullivan:* I want to note from the chat that this is consistent with the 3-d interpolator. We currently use the interpolation data and not the observed data.
- ***Response:*** *Becky Monahan:* I think that is one of the things we are concerned about with the 3-d interpolator. We didn't have high frequency data, but now we do. If you're only collecting fixed station samples, you would need to interpret between those samples. That's something that happened overtime with the 3-d interpolator. Now, with the 4-d, we have real data consistently in profile. I don't know if it's appropriate to model over real results.

- ***Response:*** *Elgin Perry:* The current interpolator using inverse distance weighting. When you are using the current interpolator to make an estimate of a point that has observed data, the weight of that data is so much greater than the other points that it practically goes through the data. It's using the interpolation, but the interpolation goes directly to the data. The 4-d interpolator does not go directly to the data. You can see in the plots that the smooth line comes close to the data but does not pass directly through. USGS has their weighted regressions on time, discharge and season (WRTDS), which does an estimation of loads coming across the fall line. They started by doing a smooth regression curve, like we do, but they later progressed to WRTDS with Kalman filter smoothing, with does nudge the smooth curve to the data. That's what Rebecca was talking about earlier. We could explore that, but we'd like to get the smooth model working before. I don't see a problem with doing the interpolation and then adding in the observed data. I assume eventually we will be computing percentages that we pump into CFD curves. I don't see a problem with throwing the observed data in with the interpolation data to compute those percentages. I haven't thought about that until it was brought up today, but I think we could put that on the table and see if we could get that to work.
- ***Comment (from chat):*** *Jon Harcum:* CB4MH is roughly a 69 km x 24 km grid (spatially)

Comment: *Tish Robertson:* This is a good discussion. The 3-d interpolator is different from the 4-d interpolator because it is completely dependent on the input or assessment data. There is no model without that data. For the 4-d interpolator, I understand that the model will be independent of the assessment data. When we are ready to use it, we could perform an assessment using a dataset that will inform or enhance that model. If that understanding is correct, then the states would benefit from having a sensitivity analysis to see how much the interpolator results change from each three-year rolling window. I want to see what those look like. I want to know how sensitive the model would be to changes in the underlying data. If it doesn't change that much, then maybe we don't have to have an assessment data set. It would cause us to have a lot of questions if that's the case.

- ***Response:*** *Rebecca Murphy:* I like that suggestion, if I understand you correctly. I think that's something we could do as we test it. We could test three-year periods at a time and see how they're different. I would imagine they are going to differ in those key parts of the Bay, where sometimes we meet the criteria and other times we don't. If we don't, is it because when we add in true high frequency variability, there's too much low oxygen or is it not accurate? That's a good suggestion.

Q: *Matt Stover:* Going back to Elgin's earlier explanation on the variograms on slide 15. The curve levels off at a greater distance, 3 and 4 km. Semi variance levels off meaning it's not getting more variable as you get further between the stations. Is that the correct way to interpret that?

- ***A:*** *Rebecca Murphy:* For the top variogram, once you reach 3km, two oxygen observations are not correlated or there is very little correlation. They're as correlated as something 4km apart. It's no longer a function of space.
- ***Q:*** *Matt Stover:* For the 1000m value, you have a semi variance value of 0.7 in the Tangier Sound and 0.6 in the Rappahannock. So, does the Rappahannock have less variance at the same distance?
- ***A:*** *Rebecca Murphy:* It has less spatial autocorrelation, which is consistent with what is on the map. The Tangier Sound has some of the most spatial autocorrelation.
- ***Q:*** *Matt Stover:* That means that two values at different locations are most similar most of the time.
- ***A:*** *Rebecca Murphy:* Yes.

- ***Q:*** *Matt Stover:* Is there a way we arrived at 1km?
- ***A:*** *Elgin Perry:* It was the existing interpolator grid.
- ***A:*** *Rebecca Murphy:* As we talked about in the last meeting, the interpolator grid does vary. Some places it's 500m and sometimes it's 100m. As Elgin mentioned, if it's closer, the 1km correlation would be adjusted to the shorter distance.
- ***Response:*** *Jon Harcum:* The 1000m distance was because it was the maximum grid distance in the existing interpolator. As Rebecca was saying, if it's 500m, we would take the square root. That is part of the reason we selected an exponential variogram. That's one of the mathematical features that lets us scale from 1000m to 500m to 250m. For those that do kriging, you may have heard of other model fits that include nugget effect and spherical models. The nice thing about using exponential variogram, which you can see in the datasets shown here (slide 15), is that it fits the data well and gives us the ability to square 1000m down to handle stations that are closer together.

Q: *Kaylyn Gootman:* I was thinking about the last discussion and have a question for Becky and Tish. I know there is a decision on how to use the data observations after the interpolating surface is calculated and how that would differ from the 3-d interpolator. From your perspectives, would it be useful to have the ability to make those comparisons? Like pulling out the observed data and comparing that to what the interpolated surface looks like. I am thinking about the integrated water quality reports and having an opportunity to increase usefulness of partnership tools.

- ***A:*** *Becky Monahan:* I think that would be really useful. That's one of the things we were looking for with Fishing Bay and having your help in analyzing it. We want to compare what the 4-d interpolator said with all the data along with just looking at the data. Obviously, it would be different because there would be interpolated values versus gaps in the data. I know each segment can be different, but it would help us understand the benefits and what it is doing in between.
- ***A:*** *Tish Robertson:* I think Becky is right on point. The monitoring effort is always going to be a big question for us. We are in an environment where monitoring funds may not always be around. Is it worth our while to do enhanced monitoring in some areas? Would it benefit the interpolator? If you get different results based on how much data is thrown in it, then we'd know it's important to do enhanced monitoring, but if there isn't much of a difference, then we can do what we've always done. We need to know how sensitive the model is to changes in the underlying dataset. That's how I would frame it.
- ***Response:*** *Kaylyn Gootman:* I appreciate that, Tish. It made me think of a future use of interpolator and observed data outputs guiding where we should direct or implement enhanced monitoring. This could be a tool to help make informed decisions and we can think about that as we move forward.

III. Addressing Stakeholders Questions

Lead: ***Breck Sullivan*** *(USGS)*

The group ran out of time and did not get to this agenda item.

IV. Adjourn

Next Meeting: ***January 26, 2026***

Attendees:

- Allison Welch, CRC
- Amanda Shaver, VA DEQ
- Andrew Keppel, MD DNR
- Angie Wei, UMCES
- Becky Monahan, MDE
- Breck Sullivan, USGS
- Efeturi Oghenekaro, DOEE
- Elgin Perry, Independent Statistician

- Jim Hagy, EPA
- Jon Harcum, TetraTech
- Joseph Morina, VA DEQ
- Kaylyn Gootman, EPA
- Leah Ettema, EPA
- Marjy Friedrichs, VIMS
- Mark Trice, MD DNR
- Matthew Stover, MDE
- Melinda Cutler, MDE
- Peter Tango, USGS
- Rebecca Murphy, UMCES
- Tish Robertson, VA DEQ