Bay Oxygen Research Group (BORG) Meeting

November 17th, 2025
12:00 PM – 1:30 PM
Visit the meeting webpage for meeting materials and additional information.

Purpose: This is the November 2025 Bay Oxygen Research Group Meeting. During this meeting, Rebecca Murphy (UMCES) presented the development team's proposal for thinning 5 segments in the 4D interpolator to address computer capability issues. The group discussed this and decided to move forward with the thinning of these five segments. They also wanted to look into thinning all of the segments with a 50m grid resolution. The development team will be looking at how that would impact the findings for those segments and bring the results back to the group. Next, Rebecca presented the sub-sampling of data in the 4D interpolator and how oversampling can impact the interpolator results. The group discussed this and will revisit the topic at a later meeting.

Minutes

I. Welcome, Introductions & Announcements
Lead: ***Breck Sullivan*** *(U.S. Geological Survey, USGS)*

II. Interpolator Grid: Proposed thinning 5 segments
Lead: ***Rebecca Murphy*** *(University of Maryland Center for Environmental Science, UMCES)*

Rebecca presented on the interpolator grid, which is in the existing interpolator. A goal for the 4D interpolator is to be able to run the interpolator on an average 32 GB laptop. When testing the 4D interpolator on a 32 GB laptop with the current grid, five segments could not run reliably. The proposed solution for this is to halve the resolution of the grids for these five segments. With this change, the analyses run on a 32 GB laptop, and the run time decreases. The interpolation results were compared between the original and thinned grids. There seemed to be little to no change in the results. Rebecca and her team wanted to know if the members of the BORG team were ok with the decision of thinning the grids for the five segments.

Decisions:

- The five segments discussed will be thinned, so the interpolator is able to run on 32 GB laptops and to improve runtime.
- The group will revisit the decision to thin the rest of the 50m resolution segments at a following meeting.

Actions:

1. The Development Team will run the simulations on the remaining 50m resolution segments to test whether they should also be thinned to 100m resolution. These results will be shared with the group where they will decide whether those segments should be thinned.

Discussion Notes:
Q: *Tish Robertson:* For example, if the Pamunkey tidal fresh has 1000 centroids that you develop an estimate for, you would go down to 500?

- ***A:*** *Rebecca Murphy:* Yes, that's right. Every other one would be removed. Instead of the centroids being 50 meters apart, they'd be 100 meters apart.
- ***Response:*** *Elgin Perry:* We are taking every other grid line in latitude and longitude so 1000 centroids would reduce to 250. It's more like being quartered than halved.
- ***Response:*** *Rebecca Murphy:* When we looked at the maps, we were pleased. For example, on slide 4, the green map shows the centroids of the 50m grid and then red is if we did the thinning and now have a 100m grid. I thought this would be problematic because we would be missing different parts of the river, but we found that it is still covered pretty well.

Q: *Amanda Shaver:* I noticed some other segments had the 50m grid cell size. How did you decide on the five that you are changing?

- ***A:*** *Rebecca Murphy:* Yes, there are 15 segments with the 50m grid size. The additional ones didn't cause any runtime trouble, so we figured that we'd keep them the way they are. We could change them, but they aren't causing any problems. When we fit these in the interpolation, there is one step that requires a matrix of cells that covers it completely. That would be like the white space in the maps on slide 4. It becomes much bigger than the area of the water. This is to do the variance covariance matrix in depth. We need to populate this entire matrix that becomes much larger than this region. If the tributary segment is curvy and angled like this one, the matrix has to be large, which can cause runtime problems. The smaller segments or segments that are oriented better don't cause as many runtime problems. If you think we don't need the 50m grid, we could change the others. We have asked around to figure out why these grids were so fine originally when this grid was created. If anyone knows the history of this grid, please let us know.

Q: *Tish Robertson:* I am curious about why the Choptank results are more variable than the Pamunkey?

- ***A:*** *Rebecca Murphy:* I think the Choptank simulations are going to get tighter. It's kind of tied to the next presentation. We've been making adjustments to the mean mid-day fit of the model and where the limits are. There is only one station close to the Choptank tidal fresh. Our variability should be tighter, but it was extrapolating into the further extent of the region.
- ***A:*** *Elgin Perry:* My current hunch is that the Pamunkey has more stations and more data, so we are getting a tighter prediction. When we get our model finalized, we will do more experiments to verify this.
- ***Q:*** *Tish Robertson:* Are you talking about the continuous data? I know we have a lot of continuous data.
- ***A:*** *Elgin Perry:* You have more fixed stations in the Pamunkey than the Choptank. I haven't compared how much ConMon data is in the two systems.
- ***Response:*** *Tish Robertson:* The Chesapeake Bay National Estuarine Research Reserve (CBNERR) ConMon is in the Pamunkey tidal fresh and that is a truly continuous station. It

was my understanding that we only had one long-term station. We may have some Virginia Department of Environmental Quality (VA DEQ) stations too.

- ***Response:*** *Amanda Shaver:* I think you are right, Tish. I think there is only one station in the Pamunkey.
- ***Response:*** *Elgin Perry:* There is only one station in the tidal fresh, but when we fit our models, we use boundary segments and therefore there are more stations in our total interpolation region, which is the target segment and the boundary segments. It is just my speculation.
- ***Response:*** *Rebecca Murphy:* We realize the Choptank tidal fresh is a good place to look at because the fixed station is right outside of it in the oligohaline. I assume sometimes there are other stations in the tidal fresh but this year there wasn't. That's a good place for us to test and make sure the results are doing what we expect when there isn't much data.

Comment (from chat): *Marjy Friedrichs:* This sounds like a great solution to the issue to me! Thanks for all your great work on this.

Comment: *Tish Robertson:* I don't have any concerns with the two segments that have been presented on. Is it fair to treat those differently than the western branch of the Elizabeth? If we feel like there is a good rationale for singling those out, then I'm ok with it. It seems like there might have been some basis for having the grid cells the way they were. I don't know if we have documentation on that.

- ***Comment:*** *Amanda Shaver:* If it's just a runtime issue, then we should consider treating them all consistently.
- ***Q:*** *Breck Sullivan:* By that, do you mean you are ok with having a longer runtime to be consistent with the others or that all of the 50m grids should be thinned?
- ***A:*** *Amanda Shaver:* Either way. We should address them all or address none of them. Runtime is a hard rationale for me in changing it.
- ***Response:*** *Elgin Perry:* This was not only a runtime issue, but also a hardware issue. If we stick with the current resolution, we'll need to limit the use of the interpolator to a workstation type unit, rather than people being able to use it on a laptop.
- ***Response:*** *Jon Harcum:* Yes, in the testing, I was getting different results on the laptop. The 32 GB limits us in the process. We would need to thin these five segments to run it on a 32 GB laptop.

Q from chat: Marjy Friedrichs: Elgin and Rebecca - is there any problem that you see with changing all the 50m resolution segments to 100m segments?

Q: *Matt Stover:* Are there any downsides to doing this for all of the segments? Is the data or interpolation method specific or precise enough to make use of the small 50x50 segments? There is always the desire to make things as precise as possible. Does our data support that in the first place?

- ***A:*** *Rebecca Murphy:* That question relates to what Marjy asked about changing all of the 50m resolution to 100m. We were sticking with the existing grid in an attempt to change as few things as possible, for you. We were scratching our heads too because 50m resolution is really

fine. I'd have to look, but it's hard to say that it's justified. I would be a little worried about the Anacostia.

- ***Response:*** *Matt Stover:* Western branch tidal fresh is also really tiny, so it may be needed there. I can't imagine that the bathometry we are using has that high of a resolution.
- ***Response:*** *Rebecca Murphy:* We do have a finer bathometry. I believe we are using a one-meter grid, but not all of the data is good, and it has its weaknesses. We did not go through and examine all of the 50m ones. We figured we wouldn't want to go too far with changes. We also couldn't figure out why the grid was so fine to begin with. Tish, we thought you might know about that. Richard Tian was not familiar.
- ***Response:*** *Tish Robertson:* They've always been the same since I started using it.
- ***Response from chat:*** *Marjy Friedrichs:* Those are exactly my concerns, and why I would support getting rid of all (or almost all) of the 50 m segments. But at least these five!

Q: *Tish Robertson:* With the resolution, are we making a statement about what constitutes an independent sample? If we go with the 50m grid cell in certain segments, are we saying that within that area you would want the data collected within that radius together and not evaluate the station separately? If that's the logic we apply to some segments, why wouldn't we apply that everywhere?

- ***A:*** *Elgin Perry:* The covariance functions we are using to describe how much dependence there is between one cell and its neighbor is a function of distance. When we shift from a 50m to 100m grid, the way we are computing the covariance among grid cells is going to make slightly more independence in the 100m grid cells. Making the adjustment for distance is built into our algorithm.
- ***Response:*** *Tish Robertson:* I was thinking that it would work like the current interpolator. You have your input data; it will aggregate the data that falls within the same cruise, and then it applies the Inverse Distance Weighting (IDW) to that aggregation. You're saying it's not the same process.
- ***Response:*** *Elgin Perry:* Yes, it's different from IDW.

Comment: *Breck Sullivan:* I have heard that we want to thin these five grids to address the ability to run the interpolator on a laptop and in an appropriate amount of time. From the discussion, we also want to look at thinning the other 50m resolution segments.

- ***Response:*** *Matt Stover:* We probably want to see similar graphs for the other ones before making that decision, so we don't have any unexpected consequences, especially those smaller tributaries.
- ***Response:*** *Elgin Perry:* I agree with you, Matt. That isn't a very big lift once we get the original simulations done. I will put that down as an action item.
- ***Comment:*** *Breck Sullivan:* We will postpone that decision and will follow up with those graphs.

III. **Sub-sampling Continuous Monitoring (ConMon) and Data Flow Data**
Lead: ***Rebecca Murphy*** *(UMCES)*

This presentation on sub-sampling helped show how all of the data is being used for the interpolator. All of the Bay's dissolved oxygen data is being used. There are four parts to the interpolation method. This presentation will focus specifically on the mean mid-day space-and-time interpolation. Previously, the development team fit the mid-day smooth functions using a daily sub-sample of high frequency data, but feedback from the group led them to try to fit the mid-day interpolation with all available high-frequency data. In doing this, they found that the high-frequency data of a single ConMon station could disproportionately affect the interpolation results at nearby stations. This is known as the "imbalance of classes" in data science literature. To remedy this, they are suggesting sub-sampling to hourly frequency, rather than 15 minutes, and every 500m for DataFlow. They may also need to make changes within the smoothing method. Rebecca showed an example of this issue, which shows how a ConMon station with a lot of data impacts interpolator findings. Then, Rebecca showed her findings from sub-sampling the 15-minute data from the previous example to 1 hour. From this, there was little to no change in the distribution of the data, meaning they are not losing anything that would change the criteria assessment conclusions. These findings were confirmed by another example. Lastly, Rebecca touched on how Dataflow would be sub-sampled, which would likely be by distance.

Actions:

1. The group will return to this conversation at a later meeting.

Discussion Notes:

Q: *Matt Stover:* The graph in the middle of slide 8 uses the ConMon data to establish what the value would be at noon?

- ***A:*** *Rebecca Murphy:* Yes, 11AM.
- ***Q:*** *Matt Stover:* In the graph on the right, the red dots are the sample results from that station, CB3.1. When you are interpolating, you aren't looking at it segment by segment, but you're including a bunch of segments together?
- ***A:*** *Rebecca Murphy:* Yes, they are all together.
- ***Q:*** *Matt Stover:* Is the data from CB3.1 influencing the interpolation in the Bush River above station WT1.1?
- ***A:*** *Rebecca Murphy:* Yes, it could be. It's included in the same smoothing function that is fit to the data. CB3.1 and XJG are fit in these smoothing functions that include relationships of oxygen with depth, bottom depth, latitude, longitude, day-of-year, and time-of-day. Because the two places have drastically different bottom depths, there shouldn't be too much influence. We're seeing influence in the opposite direction, which is an artifact that we don't want. In an ideal world, they could influence each other based on these spatial and temporal patterns.
- ***Response:*** *Matt Stover:* You're right. It should be in the opposite direction. It seems counter-intuitive to me that the 28,000 results at XJG would have that much of an impact on CB31 since they are so far apart and different. Also, there is a fixed station between them.
- ***Response:*** *Rebecca Murphy:* This is at the same depth. It's the 0.5m prediction. They have that in common, which is why that's allowed to happen, but we want to limit it because it doesn't seem appropriate.

Q: *Jim Hagy:* Can you remind us why we decided not to use one number per day around 11AM? That's making the most sense to me. This massive imbalance is creating a problem.

- ***A:*** *Rebecca Murphy:* Yes, this problem didn't exist when we pulled one value per day. The reason is that the data exists, and the BORG team has told us that we want to use all of the

data if it is available. We thought it would be ok because the smooth function would take care of it.

- ***Response:*** *Jim Hagy:* You're introducing information that is relevant to the subsequent step of the process where you consider different times a day. I see it as although it is data, it is contaminating the objective of estimating DO (dissolved oxygen) at 11AM with data from 4PM and midnight. This is the only segment that has any other time besides 11AM. We want to use all the data, but I'm unsurprised that adding 28,000 observations versus 300 creates huge statistical problems that we're asking the generalized additive models (GAMs) to overcome.
- ***Response:*** *Elgin Perry:* When I first encountered this problem, my thinking was just like yours. When the sampling is too rapid in time, it represents sampling the same unit over and over rather than getting independent observations. I also thought we should sub-sample for more independence. Then, Rebecca ran it with hourly data. It did seem to improve that and add benefit to the model. We need to be clear about the fact that for autocorrelation to be taken into account when fitting the model, we'd have to use a different tool called the generalized additive mixed effect models (GAMM). It's a maximum likelihood algorithm that uses iteration to get its optimal solution. We tried some experiments early on and found that in a lot of cases it doesn't converge. It didn't seem like it was feasible to use that tool in a package that would be a production tool. The GAM package we're using does assume independence of the observations. I think that's one of the things that hurts us when we try to put 15-minute data into the model. One way to de-emphasize these is to sub-sample them or do a weighted least squares fit, which could use the hourly data. Those are all options that are proposed in the data science literature I am looking at. I'm sure this is something we will come back to when we have more time to research it.

Comment: *Carl Friedrichs:* At the Coastal Estuarine Research Federation (CERF) Conference, Rebecca explained to me that it could be important to use data from other times in the GAM, like to calculate daily values. I don't like the idea of throwing data out because it is causing issues. Statistically, you're adding data that is locally autocorrelated in time, which could support the argument for dropping to hourly. I'm not sure what degree the statistics of autocorrelation take care of that. I think you're oversampling, and there are statistical reasons not to use data that is oversampled in one place. That could sound better than throwing out data, but instead a statistical oversampling argument. The spatial correlation that prevents data from far away being used could fix the issue along with taking care of the oversampling. I'm fine with taking that data away, but we should attribute it to the statistical issues that arise from oversampling in autocorrelated models.

Comment: *Matt Stover:* The fixed stations aren't all sampled at 11AM. There is a bit of modeling just to get at an 11AM value. I don't think it would be the best thing to put so much bias towards the fixed station. I don't see it as a contamination of the assessment by ConMon data but more so a problem with how we're geographically interpolating. CB3.1 is a long way from XJG. It raises the question of whether they should be communicating when we're assessing DO. Our DO criteria have high frequency components in them. Using monthly or twice monthly data doesn't allow us to assess our criteria. Maybe this argues for doing a simpler station by station assessment, like Maryland and Virginia have been talking about. It seems counterintuitive to do it that way. We rarely complain about having too much data. It seems like a waste to not use it in a way that is meaningful to assessing DO criteria, at least in the locality where the data was monitored.

- ***Response:*** *Rebecca Murphy:* You mentioned that fixed station data has a time associated, and we do include time. The time that the fixed station data was collected is linked to the time-of-day term in the smoothing function. We picked 11AM because Jon did a ton of analysis of

fixed data time and it centers around 11AM across the entire data record. When it varies, we get that change in the time-of-day which is accounted for in the tool.

Q from chat: *Jim Hagy:* Could you de-weight the data, so you use it all, but the weights are something like 300/28000?

- ***Comment from chat:*** *Rebecca Murphy:* Thanks for that idea, Jim.
- ***Response from chat:*** *Jim Hagy:* You could also sub-sample the ConMon data so that it resolves time of data effect within the range that discrete data vary, which is to say something like 8 AM to 4 PM. But then not consider data from 4 PM to 8 AM. I imagine that the diel cycle at CB3.1 does not look very much like the diel cycle at XJG7035, which is likely more dynamic.

Comment: *Kaylyn Gootman:* This is a good discussion. It's great to see the visuals. I wanted to bring up the continuous monitoring question of using so many pieces of information. This is still cutting edge and at the forefront of monitoring. We are thinking about this on the nontidal side with rectifying continuous monitoring. I am hesitant to truncate. I think we have the unique opportunity to use that data in sub-sampling and de-weighting in the way that makes the most sense. There is so much going on in these non-sampling hours. Understanding what is happening at those off hours is really valuable and puts emphasis on the need to have continuous monitoring. I see value in including those other hours however we can.

Q: *Tish Robertson:* When you say sub-sample is an hourly average used or are you randomly selecting a data point?

- ***A:*** *Rebecca Murphy:* I took one sample every hour. I did not average.

Q: *Tish Robertson:* With the dataflow at 500m, what does that mean if we had fixed stations within 500m, if they were closer than 500m?

- ***A:*** *Rebecca Murphy:* Those are different situations. If there were two stations within 500m, I would still think they should stay as independent units where we keep their data. This would just be for dataflow data because the data is so spatially and temporally dense.

Q: *Tish Robertson:* Would it be like the sub-sampling with the ConMons but through space? Where would you sub-sample within 500m increments?

- ***A:*** *Rebecca Murphy:* We will select a sample every 500m starting at the start point.

Q: *Tish Robertson:* What's the benefit of sub-sampling over calculating an average?

- ***A:*** *Rebecca Murphy:* The average would change the nature of the data point. Because we are doing discrete data points, we wouldn't want some of them to be averages, so it would be a different type of data.
- ***Q:*** *Tish Robertson:* You would be adding more variability or error if we averaged it?
- ***A:*** *Rebecca Murphy:* They would not have the same distribution. In the statistical models, the data should be similarly distributed. The distribution would be lower on the averages. We are trying to stick with the observed variability of the data.
- ***Comment from chat:*** *Jim Hagy:* The error distribution would be different... lower.

Comment: *Carl Friedrichs:* I used this paper when we had oversampled tidal dataflow data. We were sampling water quality in an aquaculture farm. We were doing statistics on how many effective samples it was. There is a simple method where you can get the equivalent independent observation by the autocorrelation length. It's a simplistic argument and lets you choose what

effective number of samples to use to get a more realistic number of observations. It worked well for us and might be worth looking at.

- Zieba A, Ramza P. Standard deviation of the mean of autocorrelated observations estimated with the use of the autocorrelation function estimated from the data. Metrol Meas Syst. 2011;XVIII: 529–542. https://yadda.icm.edu.pl/baztech/element/bwmeta1.element.baztech-article-BSW1-0087-0002

Comment: *Breck Sullivan:* We will bring this item back for follow-up discussions along with the graphs for the interpolation grid of the 50m segments. I am hearing that we want to make sure we are not just throwing data out but are able to represent it across the segments. In an ideal world, we would have monitoring data everywhere, but that isn't feasible. We need to make sure we are able to estimate DO where there isn't data. We will bring these conversations back. I know we weren't able to fit in the open discussion period, so maybe we could do two hours for a future meeting or maybe more frequent meetings. If you are interested in the information we were going to share to address that question, I would recommend going to the webpage and looking at Jon Harcum's presentation, specifically the figures on slides 2 and 4, and Rebecca Murphy's presentation. If you have questions that come up, feel free to send them our way.

IV. Addressing Stakeholders Questions

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

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

V. Adjourn

Next Meeting: February 2026 (possibly earlier)

Attendees:

- Allison Welch, Chesapeake Research Consortium (CRC)
- Amanda Shaver, VA DEQ
- Andrew Keppel, Maryland Department of Natural Resources (MD DNR)
- Angie Wei, UMCES
- Becky Monahan, Maryland Department of the Environment (MDE)
- Breck Sullivan, USGS
- Carl Friedrichs, Virginia Institute of Marine Sciences (VIMS)
- Cindy Johnson, VA DEQ
- Elgin Perry, Independent Statistician
- Guido Yactayo, MDE
- Jay Lazar, National Oceanic and Atmospheric Association (NOAA)
- Jim Hagy, U.S. Environmental Protection Agency (EPA)
- Joseph Morina, VA DEQ
- Jon Harcum, TetraTech
- Kaylyn Gootman, EPA
- Leah Ettema, EPA
- Marjy Friedrichs, VIMS
- Matt Stover, MDE
- Melinda Cutler, MDE
- Rebecca Murphy, UMCES
- Richard Tian, UMCES
- Tish Robertson, VA DEQ