OPEN NAVIGATION SURFACE WORKING GROUP

MEETING SUMMARY

2017-03-21
U.S. Hydrographic Conference

DRAFT VERSION
2017-03-30

1 Introduction

This document details the result of the meeting held at 1530CST on 2017-03-21 to continue development for a new release of the library. The meeting was held in conjunction with the U.S. Hydrographic Conference 2016, using facilities kindly arranged by Melissa Wood of the Hydrographic Society of America. The summary of all meetings and teleconferences of the Open Navigation Surface Working Group (ONSWG) can be obtained from the project's web-site, http://www.opennavsurf.org. For a list of participants, see section 4.

In the following, names of people with action items are shown in **BOLD SMALL CAPS**; expected deadline release dates are shown in red. Sizes of variables are indicated by 'U' for unsigned, 'S' for signed, 'F' for floating-point, and a size in bits (e.g., U8 is an eight bit unsigned integer, F64 is a 64-bit (double precision) floating-point number). Data sizes are given in bytes (B) with the usual convention that the SI multipliers are taken to mean multiples of 2^{10}B (i.e., 1kB = 2^{10}B = 1024B). The acronym 'CR' means 'Candidate Release' (i.e., a release of the library for comments) and 'FR' means 'Full Release' (i.e., release V1.7.0 of the library).

2 Summary of Discussion

2.1 Requirement for Nominal Depth Layers

In the current release of the library, it is possible to have Nominal Depth stored in the elevation layer of the BAG file instead of real depths; a metadata record is used to indicate to the user which type of data is present. At the ONSWG meeting held at the Canadian Hydrographic Conference in 2016, Naval Oceanographic Office personnel indicated that this requirement is no longer essential for their operations, and **Ladner** indicated that a grace period would be beneficial in realigning the production process at NAVO to allow for this.

The group discussed this requirement, and agreed that removing Nominal Depth as an option for the primary Elevation Layer was still the right thing to do. **Ladner** requested, however, that this be approached by deprecating this use in the next release (1.7.0) of the library, and dropping it on the following release. Agreed *nem. con.* [Action: **CALDER**].

2.2 Categorical Description of Uncertainty in Metadata

At the last ONSWG meeting, the group discussed the need for a standard definition of uncertainty to be associated with the "productUncertainty" metadata tag. After polling the development mailing list, **Calder** reported that there were at least 12 difference descriptions of the uncertainty, although after some simplification this could be reduced to three basic alternatives, so long as there were agreed axiomatic descriptions for treatment of designated soundings, and the option for a minimum uncertainty to be specified in the metadata. The options currently in use are, specifically:

- CARIS standard uncertainty. Larger of CUBE's uncertainty output and the standard deviation of all soundings in a bin around the node, reported at 95% CI assuming Gaussian statistics.
- Leidos standard uncertainty. Larger of CUBE's uncertainty output and the average vertical uncertainty of all of the soundings that make up the primary hypothesis in CUBE, reported at 95% CI assuming Gaussian statistics.
- Generalized standard uncertainty. Uncertainty derived from geostatistical estimation or other methods, described with a statement in the metadata.

The group discussed the options for harmonizing, or at least distinguishing, these options. The consensus conclusion was that it would make more sense to come to a common definition of how to report the uncertainty in the product, rather than trying to provide several different alternatives. The participants felt, however, that any decision to change the reporting was something that required further consultation within their respective organizations before adoption. **Calder** therefore agreed to make a specific proposal to the development e-mail list for this change, with a tentative agreement that the specification should be that "productUncertainty" mean the larger of CUBE's standard uncertainty output and the standard deviation of the soundings used to construct the primary hypothesis. The proposal would also necessarily have an additional metadata tag for "generalizedUncertainty" (as a catch-all with a description in the metadata), and a

requirement for full and explicit definition of the meaning of these uncertainty tags in the documentation for V1.7.0. [Action: CALDER].

2.3   Inconsistent Georeferencing Between Vendors

Since the last meeting, **Paton** reported via the development e-mail list that there appeared to be an error in how georeferencing information was being interpreted by different vendors. Specifically, this is an issue for some vendors that rely on the use of the GDAL library to read the data, which bypasses the georeferencing information in the BAG library itself. After investigation, it appears that the issue is related to providing more information than is required to specify the georeferencing; specifically, providing a southwest and northeast bounding corner, plus a sample spacing, plus the number of rows/columns. The BAG library uses one interpretation of this data to compute the georeferencing, ignoring the northeast corner (which might be slightly misaligned); other readers do not. Paton agreed to provide the detailed analysis of this matter to the development e-mail list, noting that other issues may have occurred in the past (e.g., mistakenly moving the grid by a half-cell, and then correcting for this on readback) [Action: PATON].

The group discussed the issue, and concluded that it would be possible to resolve the issue by adding a check to the BAG creation call that verified the georeferencing information, failing the call if invalid; adding a check for bad georeferencing data on read, with automatic correction, in order to support older BAG files; and to increase the number of decimal places being reported for the coverage metadata to have more decimal places (currently only two places are used, which is not adequate in geographic coordinates). [Action: PATON].

2.4   Documentation for Georeferencing Conventions

Van Duzee reported that there appeared to be an issue with the Well-Known Text (WKT) used since V1.5 to provide georeferencing information for BAG files. Specifically, it is possible to define a datum transformation within the WKT through a series of parameters, but there is no way to mandate how those parameters are used. Unfortunately, there are multiple ways of interpreting the parameters, and no way to tell which method is intended. The BAG library currently uses EPSG 9606 (position vector) rather than the alternate EPSG 9607 (coordinate frame rotation); vendors that do not use the BAG library to read the data therefore have the potential to misinterpret the intent of anyone using the library to write the WKT. This is believed to be an active issue with the GDAL library implementation.

The group agreed that this could be resolved simply by defining the intended usage axiomatically as part of the documentation for V1.7.0 of the library. The most likely usage is EPSG 9606, but the group agreed that it was useful to poll the development e-mail list to ensure that other usages were not active before moving ahead. [Action: CALDER].

2.5   Potential for Active Tiling in BAG Files

The idea of allowing the BAG library to actively tile the Elevation and Uncertainty layers (to exploit sparseness in many surveys for data volume reduction in addition to compressing the data) was discussed at the 2016 ONSWG meeting. The group discussed the proposal again, but concluded that limited developer resources for this were better spent on construction of the V2.0 API (see Section 2.6), and that tiling should therefore be considered as part of the specification for V2.0 BAGs, rather than attempting to patch this functionality into the increasingly dated V1.0 API.

2.6   Proposal for Revised C++ API

The development group have previous discussed the need for a new API for the library, specifically to take advantage of modern C++ design techniques, and to simplify access patterns for the various layers of BAG files. Prior work on this by CARIS is available as the 'v2.0_cpp_api' branch of the Git repository, although **van Duzee** noted that the implementation is at present incomplete.

The group agreed that a new API is still required, and further agreed that developers should assume that the API will be purely C++, and that C++11 language features would be allowed (i.e., the compiler versions used by all of the primary vendors would support C++11 features). The group also agreed that there should be no constraint that V2.x API should use legacy code from V1.x, or necessarily support the same access patterns: V2.x API is expected to be a complete re-write of the library, without constraint. However, it is a requirement that V2.x API versions of the library will be able to read V1.x series BAG files, although not write them. The binary structure of the BAG file is not expected to change, except to support features

new to V2.x, such as tiling. However, code using V1.x-series libraries are not expected to be able to read BAG files written with V2.x-series libraries.

Development of V2.x will also provide the opportunity to update to the latest HDF5 API. Current version of the library use an antiquated version of the HDF5 API in backwards compatibility mode; although still currently supported, this mode will eventually be deprecated, and the group agreed that the BAG library should take the opportunity to update. Paton agreed to investigate and report how the new HDF5 API will integrate with the BAG library requirements. [Action: **PATON**].

In discussion, the group noted that it would be beneficial to have as many different types of BAG files (from different vendors and version of the software) to provide regression testing opportunities. Calder agreed to support these on the project website, but asked for vendors to forward the files. [Action: **CALDER, ALL**].

2.7 Storage of Backscatter Data

Masetti posed the question of whether it would be of interest to include backscatter information in the BAG file, or to have a variant of BAG that allowed for this. The goal is to provide a standard product for backscatter, which is currently not available. The argument was that many backscatter processing methods rely on the use of bathymetry information, and therefore that having them together in one place would potentially be useful.

The group agreed that including backscatter directly in a BAG file would be problematic, but that a sister-project using the same technologies (and potentially sharing the library) might be more appropriate. Both **QPS** and the **Naval Oceanographic Office** expressed interest in developing such a project. The ONSWG encourages development of this project, and are willing to assist if possible; no action is required at this time, however.

2.8 Vertical Referencing for Radar-Derived Bathymetric Data

Riley had previously posed the question of which vertical datum should be specified for data which extends over a significant period, but does not correct for this; the specific example was bathymetry derived from X-band radar observations, where significant temporal averaging is used to generate the bathymetric estimates, which makes it difficult to correct for tidal effects, and therefore reduce to a datum.

The group briefly discussed the proposal, and agreed that it would be appropriate to recommend a specific EPSG code, but were unclear what it should be. After the meeting, **Riley** indicated that the group generating this data had concluded that they could report to Mean Sea Level, obviating the need to select a EPSG code. No action is therefore required at this time.

2.9 Documentation System for the Project

The group discussed the state of the documentation for the project, and concluded that using a Word file for the File Specification Document was limiting the ability of the project to maintain the documentation in an appropriately flexible manner. **Paton** suggested that it might be possible to use Atlassian Confluence to implement the documentation, noting that **QPS** uses this internally for their documentation, and find it useful. (**CARIS** also indicated that they use Confluence internally, and are happy with the system.) There is a potential concern about costs associated with the use of this package for a non-funded project, but it is likely that free licenses may be available for such projects. **Beaudoin** agreed to investigate this matter, and report to the group. [Action: **BEAUDOIN**].

3   Summary of Action Items and Dates

The following actions and dates were agreed:

Person	Actions(s)	Section	Date
Calder	Add documentation to 1.7.0 to deprecate use of Nominal Depth for the primary Elevation Layer	2.1	2017-07-21
Calder	Proposal to dev. list for explicit definition of Uncertainty Layer tags	2.2	2017-04-04
Paton	Provide georeferencing analysis to development e-mail list	2.3	2017-04-04

Paton	Implement changes to georeferencing handling for consistency on creation, and fix-up on read.	2.3	2017-07-21
Calder	Poll development e-mail list on preference for interpretation of datum transform in WKT	2.4	2017-04-04
Paton	Investigate and report on how the new HDF5 library API will support V2.x-series BAG libraries.	2.6	2017-06-01
Calder	Upload BAG file examples to website.	2.6	2017-07-21
All	Provide example BAG files to Calder for update of website.	2.6	2017-07-21
Beaudoin	Investigate requirements for using Confluence as a documentation system for the project.	2.9	2017-04-04

Dates above in red are those which would result in a significant impact on other activities were they to slip, and are therefore critical. The release dates for 1.7.0 agreed were:

- Candidate: 2017-07-21
- Full: 2017-08-04

4 Participants

Jonathan Beaudoin (QPS)
Brian Calder (CCOM/JHC)
Paul Donaldson (Leidos)
Mike van Duzee (CARIS) [Teleconference]
Burns Foster (CARIS)
Jason Infantino (Leidos)
Wade Ladner (NAVO)
Giuseppe Masetti (CCOM/JHC)
Danny Neville (QPS) [Teleconference]
Mark Paton (QPS) [Teleconference]
Matt Thompson (NAVO)
E. J. Vandenameele (NOAA)