NWP SAF

MWIPP Version 1 Test Plan

Version 1.0

13th November 2018

MWIPP Version 1 Test Plan

This documentation was developed within the context of the EUMETSAT Satellite Application Facility on Numerical Weather Prediction (NWP SAF), under the Cooperation Agreement dated 7th December 2016, between EUMETSAT and the Met Office, UK, by one or more partners within the NWP SAF. The partners in the NWP SAF are the Met Office, ECMWF, DWD and Météo France.

Version	Date	Author	Remarks
0.1	13.9.18	N Atkinson	Initial draft
0.2	25.9.18	N Atkinson	Start adding test results
1.0	13.11.18	N Atkinson	Beta test results and responses have been included

Table of Contents

1. INTRODUCTION

This document defines the test plan for Version 1 of the Microwave Imager Pre-processor (MWIPP). MWIPP is a deliverable of the NWP SAF.

As stated in the NWPSAF CDOP-3 proposal, the MWIPP will be a generic pre-processor for microwave imagers based on the existing capability of the SSMIS-PP package, developed during CDOP-1 and 2. Initially this package is intended for use with SSMIS, AMSR-2 and FY-3 MWRI. Later releases will include MWI and ICI on EPS-SG.

The primary purpose of the package is to prepare data for input to NWP assimilation systems.

NWP SAF software products are developed according to the guidelines of [AD-1]. The purpose of the testing described in this document is to ensure that the requirements of the Product Specification [AD-2] are met and that the design of [AD-3] is appropriately implemented.

1.1 Applicable documents

[AD-1] NWP SAF Development Procedures for Software Deliverables: NWPSAF-MO-SW-002
[AD-2] NWP SAF MWIPP Product Specification: NWPSAF-MO-DS-035
[AD-3] NWP SAF MWIPP Top Level Design: NWPSAF-MO-DS-034

2. TESTED ITEMS

2.1 Coding inspection

The code will be inspected by someone other than the author. The inspection will be to sign off that the code

- Is written to the guidelines of AD-2 (see below)
- Correctly implements the design of AD-3

[AD-2] recommends that each source code module should contain, in its header (or near the top of the code):

- The EUMETSAT / NWP SAF copyright statement
- The function of the module/subroutine
- Author, creation date and version number
- History of changes
- Definitions of variables
- How it is called
- Which subroutines and modules it calls
- Inputs and outputs

Additionally, as part of the coding inspection the web sites appropriate to the external libraries should be checked, to ensure that appropriate versions of the libraries are available and that licensing arrangements are acceptable.

2.2 Compilation testing

The code should compile without error and with no unexpected warnings. The tester is to follow the instructions in the Installation Guide, to run the "configure" and "make" scripts. If the required external libraries (ecCodes and hdf5) are not already installed, these must be installed first.

2.3 Module testing and Integration testing

MWIPP is a relatively small package, created by a single developer. The development approach chosen is to start with a small program performing a basic function (e.g. ingest of hdf5 files) and to build up the complexity by adding extra functionality, e.g. BUFR output, re-mapping, extra instruments, etc. – see the list of high-level requirements in the Product Specification. Therefore the Module testing and Integration testing have been combined in this Test Plan.

These tests involve the running of test cases. The purpose of these tests is to ensure that the software runs without error, that all the required functionality is present and that the results look reasonable. The final step of each test is to generate quick-look imagery showing that the output is as expected, e.g.: imagery matches up with coastlines; spatial smoothing is visible; re-mapping to the user grid has been done. For this reason, a python plotting utility is included in the deliverable. It is not mandatory for users to be able to run this utility, but it could be helpful.

In addition, BUFR output can be inspected with the *bufr_dump* tool, and hdf5 output can be inspected with *h5dump*.

The test cases (to be used both before and after release) are detailed in Table 1. These tests exercise the main functionalities of MWIPP.

Table 1: MWIPP standard test cases

Instrument	Script	Functionality tested
SSMIS	run_ssmis_upp.sh	1. BUFR ingest; averaging; BUFR output (as in SSMIS_PP) 2. BUFR ingest; averaging; map to t255_n128 user grid; hdf5 output 3. Convert averaged and unaveraged BUFR files to hdf5 and display them as a map.
AMSR-2	run_amsr2.sh	1. Standalone hdf5 to BUFR conversion for GCOM-W1 AMSR-2 2. AMSR-2 hdf5 ingest; generic hdf5 output; map to regular N320 user grid; output on user grid in BUFR and hdf5; display the original and mapped data.
GMI	run_gmi.sh	1. Standalone hdf5 to BUFR conversion for GMI 2. GMI hdf5 ingest; generic hdf5 output; averaging; map to t511_n256 user grid; output on user grid in BUFR and hdf5; display the original and mapped data.
MWRI	run_mwri.sh	1. Standalone hdf5 to BUFR conversion for FY-3C MWRI 2. MWRI hdf5 ingest; generic hdf5 output; averaging; map to t511_n256 user grid; output on user grid in BUFR and hdf5; display the original

		and mapped data.

Additionally, ingest of EUMETSAT BUFR files for GMI and AMSR-2, and re-generation of the BUFR output (dotted line in Fig 1 of the Top Level Design) should be tested.

2.4 Validation testing

Validation involves the use of independent data to confirm that the product meets the requirements. It is therefore more quantitative than the tests described in 2.3.

For SSMIS, comparisons are to be made between the MWIPP output and that of the SSMIS UPP Averaging Module. Small differences are to be expected (due to the use of fixed look-up tables in SSMIS), but differences should be small relative to typical NWP O-B values (see the data quality section of the NWPSAF monitoring web pages).

For the AMSR-2 and GMI hdf to BUFR converters, comparison can be made between the MWIPP output (generated from the hdf5) and the equivalent BUFR data sent via EUMETCast.

2.5 Portability testing

The MWIPP software is required to run on a 64-bit Linux PC, running a current operating system such as CentOS7 or RHEL7. There is no requirement to run on multiple platforms.

The choice of Fortran compiler is determined, in part, by the compilers supported by the external libraries (ecCodes and hdf5). Currently gfortran and ifort are supported.

Portability testing will therefore comprise building MWIPP using gfortran and ifort, and verifying that the test cases run correctly for each compiler.

2.6 Timing testing

Run times will be noted during the course of the Module and Integration testing, and included as information for users.

Additionally, there is a requirement that the run time for the SSMIS averaging functionality should not be significantly longer than that of SSMIS-PP. It may be slightly longer, since generic software is often slower than non-generic, and it is known that ecCodes tends to run slower than the older BUFRDC. The goal is that MWIPP should not be slower than SSMIS-PP by more than a factor ~2. If it is longer than this, justification should be provided in the Test Log.

2.7 Regression testing

If any changes are made to the software after the start of formal testing, it may be necessary to re-run some or all of the previous tests. The impact of any changes between release versions shall be carefully monitored and reported in the Test Log.

2.8 Documentation testing

The MWIPP user documents are to be reviewed by beta testers. These documents are:

- NWPSAF_MO_UD_044: MWIPP User Manual. This document combines the Scientific Description, Software Description and Installation Guide
- The Release Note (initially a draft).
- Instructions for running test cases (supplied as README files).

The beta testers should also have access to the Product Specification (AD-2) and Top Level Design (AD-3). The beta testers will provide feedback about their experiences. The developers will address any recommendations for improvement before release of the package to general users.

3. TEST CONFIGURATON

The following need to be stated in the Test Results:

- Version of ecCodes
- Version of hdf5
- Version of MWIPP (i.e. the subversion release number)
- Basic information about the system, as reported by `uname -a`
- System memory and number of processors

4. REQUIREMENTS TRACEABILITY MATRIX

Table 2 relates the requirements of the Product Specification to specific sections of the Test Plan.

Table 2: Requirements traceability matrix for MWIPP

Identifier	Requirement in Product Spec	Testing method	Test plan reference
MWIPP1	Documentation is clear, understandable and complete	Beta testing	2.8
MWIPP2	Code conforms to the requirements of NWPSAF-MO-SW-002: commented, understandable and modular	Inspection	2.1
MWIPP3	Any necessary external libraries are freely available	Inspection (e.g. examine the corresponding web sites for the external libraries)	2.1
MWIPP4	Code builds with no errors on a 64-bit Linux PC, running a current operating system such as CentOS7 or RHEL7. More than one Fortran compiler shall be tested.	Test	2.2, 2.5
MWIPP5	Ingest BUFR files for SSMIS and AMSR-2	Test	2.3
MWIPP6	Ingest native-format files for AMSR-2	Test	2.3
MWIPP7	Spatial averaging capability (replicates SSMIS_PP)	Test	2.3, 2.4
MWIPP8	Map to user-defined grid	Test	2.3
MWIPP9	Creation of BUFR output files for SSMIS and AMSR-2 (optionally GMI and MWRI)	Test	2.3, 2.4
MWIPP10	Creation of hdf5 or NetCDF4 output files	Test	2.3

MWIPP11	Run times are documented in the test log, and are comparable with SSMIS_PP when processing SSMIS.	Test and inspection	2.3, 2.6
MWIPP12	Test cases for the users exist, have clear instructions and run correctly	Beta testing	2.8

5. TEST RESULTS

The tables in the following subsections will be completed as testing proceeds.

5.1 Tests to be carried out at the Met Office

Table 3: Test setup

MWIPP version	1817 2018-09-26 15:45:24
ecCodes version	2.8.0
hdf5 version	1.8.16
Machine name	eld030
`uname -a`	Linux eld030 2.6.32-754.3.5.el6.x86_64 #1 SMP Thu Aug 9 11:56:22 EDT 2018 x86_64 x86_64 x86_64 GNU/Linux

Memory
`cat /proc/meminfo`
	`Total: 6116832 kB`
	`MemFree: 778036 kB`
	`Inactive: 1298048 kB`

Processors
`lscpu`
	`CPU(s): 8`
	`Thread(s) per core: 2`
	`Core(s) per socket: 4`
	`Socket(s): 1`

Table 4: Test results – Met Office

Category	Specific test	Performed by	Date	Result
Coding inspection	Review code	BC	09.10.18	OK, subject to minor issues detailed below
	Check external libraries and licenses	NCA	26.09.18	ecCodes 2.8.2 is the latest. Other releases available at https://confluence.ecmwf.int/display/ECC/Releases ecCodes released under Apache v2.0 license. See https://confluence.ecmwf.int/display/ECC/License http://www.apache.org/licenses/LICENSE-2.0 hdf5 1.10.3 is the latest. Other releases available at https://portal.hdfgroup.org/display/support/Download+HDF5 hdf5 license here: https://www.hdfgroup.org/licenses. There are no restrictions on use.
Compilation testing	gfortran	NCA	26.09.18	OK
	ifort	NCA	26.09.18	OK
Module testing and Integration testing	SSMIS, gfortran	NCA	26.09.18	Ran OK. Part 1 (SSMIS_PP function): 20s Part 2 (with re-mapping to t255_n128): 38s
	AMSR-2, gfortran	NCA	26.09.18	Ran OK. Part 1 (hdf5 to BUFR): 10s Part 2 (re-mapping to regular N320): 5s
	GMI, gfortran	NCA	26.09.18	Ran OK. Part 1 (hdf5 to BUFR): 0.6s Part 2 (averaging and re-mapping to t511_n256): 14s

	MWRI, gfortran	NCA	26.09.18	Ran OK. Part 1 (hdf5 to BUFR): 3.5s Part 2 (averaging and re-mapping to t511_n256): 65s
	SSMIS, ifort	NCA	26.09.18	Ran OK. Part 1 (SSMIS_PP function): 30s Part 2 (with re-mapping to t255_n128): 49s
	AMSR-2, ifort	NCA	26.09.18	Ran OK. Part 1 (hdf5 to BUFR): 10s Part 2 (re-mapping to regular N320): 5s
	GMI, ifort	NCA	26.09.18	Ran OK. Part 1 (hdf5 to BUFR): 0.5s Part 2 (averaging and re-mapping to t511_n256): 14s
	MWRI, ifort	NCA	26.09.18	Ran OK. Part 1 (hdf5 to BUFR): 3.9s Part 2 (averaging and re-mapping to t511_n256): 72s
Validation testing	Comparison with SSMIS/PP	NCA	21.09.18	See analysis in 6.1. Some small differences are noted. These are considered acceptable.
	EUMETSAT BUFR for AMSR-2	NCA	25.09.18	See analysis in 6.2. Some minor differences are noted.
	EUMETSAT BUFR for GMI	NCA	25.09.18	See analysis in 6.3. No significant differences.
Portability testing	Test cases run on a different machine (exprsatappdev01)	NCA	26.09.18	Results consistent with those given above. Gave a failure if namelist contained comments starting in the first line, therefore put a note in the User Guide about this.
Timing testing	Timings recorded	NCA	26.09.18	Standard output, with timings, is redirected to the “reference products” directory in each test case.
	Comparison with SSMIS_PP	NCA	26.09.18	Factor 4 slower than SSMIS_PP, due largely to the use of ecCodes (compared with Met Office BUFR library). Waiver requested.
Regression testing	If re-test is required, record it here.			
Documentation testing	Internal review	BC	09.10.18	See detailed comments below

Detailed comments from Met Office reviewer (BC) are given below. Developer responses in red.

Overall Summary

The code is well written and understandable. It follows correct standards for Fortran and is likely to be maintainable by someone who is not the author. The user guide is also clear. I have found a few minor corrections and also suggest a few clarifications below.

1. Review test plan

6.3 GMI. first line refers to AMSR. Suspect it is GMI Ditto Table 8. Replace with GMI? This has been corrected

2. Review user guide

- p25 "on the fly". a bit colloquial? Suggest "during execution of the program". Agreed - corrected

- better ref for MWRI p27 suggest.

Yang, H., Weng, F., Lv L., Lu N., Liu G., Bai M., Qian Q., He J., and Xu H., 2011, The FengYun-3 Microwave Radiation Imager On-Orbit Verification, IEEE TGARS, Vol 49(11). Agreed – the reference has been changed.

3. Detailed Code Review
----

read_averaging_namelist:

All elements of ChannelsToBeAveraged should be initialised to 0 before reading the namelist file
I agree it needs to be initialised. Actually I've initialised to -1 to indicate "not set"; if this is not changed by the user then the code sets default averaging for all channels. The user can enter a value "0" if no channels are to be averaged.

mwipp_latlon_mod:

I checked that the conversion to x,y,z from lat long is scientifically correct.

mwipp_averaging_mod,compute_weights:

lines 212-216 remove comments for unused write statement. Done

mwipp_averaging_mod,modify_btemps:

I verified that algorithm used is as described in user manual.

ssmis_main:

line 216 suggest adding 'regular' to run time comment. Done. Also for the other instruments.

project_onto_usergrid:
this uses ec codes routine codes_grib_find_nearest lat lon location. A status flag is returned and if non zero a nearest lat long location on the user grid cannot be found.
This means the transform to the user grid hasn't worked. The code issues a warning, but a scan lat long location is still returned. Should this not be skipped?
The ecCodes documentation is unclear as to what happens if some points in the input array are within the grid and some are not. Therefore I ran a little test on ecCodes: it seems that if just some points are bad then the return value is 0, but if all points are bad the return code is non-zero. Therefore I agree that the loop should be skipped if the return code is non-zero. I also found that

for points that can’t be mapped all the returned values are zero (index, distance, etc.) therefore have added a test to exclude any such points.

amsr2 processing:

amsr2_read_hdf: line 82. 89GHz is sampled twice spatially . Probably worth mentioning that here to explain *2. A comment has been added.

mwri processing:

mwri_read_hdf: line 116 Is setting dims(3) to 0 equivalent to passing dims(1:2) to subsequent call to h5dread_f? In effect, yes. H5dread_f will ignore any elements of dims that are zero.

4. Building and Running

Code builds ok following user guide instructions.
Test cases - I have tried GMI and MWRI tests (these are the instruments I am most familiar with). All ran ok. Inspection via the python viewer gave sensible results.
I investigated the effects of a) running with changes in sigma distance for the averaging and b) moving to a regular grid of the size comparable with the 4dVar grid here.
Results all look sensible ...N.B. I viewed channel 8 which is 89 GHZ so has small scale features over land. Easier to check averaging effects. Noted.

5. Outlook

I can imagine future releases might require more sophisticated regridding operations. But the ones here for now execute quickly.

5.2 Beta test report from ECMWF

Responses from developer are in red.

Table 5: Test setup: ECMWF

Institute	ECMWF
MWIPP version	0.1
ecCodes version	2.7.3
hdf5 version	1.10.4
Machine name	lamorak
`uname -a`	Linux lamorak 4.4.76-1-default #1 SMP Fri Jul 14 08:48:13 UTC 2017 (9a2885c) x86_64 x86_64 x86_64 GNU/Linux
Memory `cat /proc/meminfo`	MemTotal: 32868044 kB MemFree: 289736 kB MemAvailable: 25161096 kB Inactive: 13434996 kB

Processors `lscpu`	Architecture: x86_64 CPU op-mode(s): 32-bit, 64-bit Byte Order: Little Endian CPU(s): 8 On-line CPU(s) list: 0-7 Thread(s) per core: 2 Core(s) per socket: 4 Socket(s): 1 NUMA node(s): 1 Vendor ID: GenuineIntel CPU family: 6 Model: 58 Model name: Intel(R) Core(TM) i7-3770 CPU @ 3.40GHz Stepping: 9 CPU MHz: 1652.648 CPU max MHz: 3900.0000 CPU min MHz: 1600.0000 BogoMIPS: 6784.16 Virtualization: VT-x L1d cache: 32K L1i cache: 32K L2 cache: 256K L3 cache: 8192K NUMA node0 CPU(s): 0-7

Table 6: Beta test results: ECMWF

Category	Specific test	Beta tester comments
Coding inspection	Review code, at the discretion of the beta tester	
Compilation testing	Tester to select compiler	Gfortran from GNU 6.3.0
User test cases	SSMIS	Part 1: 12 s Part 2: 20 s Part 3: 5 s Part 4: 5 s BUFR files Almost all contents identical to the Met Office reference at the precision of bufr_dump but occasional differences in brightness temperature are present, usually 0.1K but sometimes as large as ~10K. Established that the set of reference results provided was not up to date. *A new set of reference results will be generated before release*. Also applies to the differences noted below. HDF file NS.h5 All brightness temperatures and other contents appear identical to reference HDF file NS_AV.h5 Patchy differences compared to reference (looking a little like mesoscale convection, but too infrequent and not in climatological locations) of at least +/- 5K.

		HDF file NS_AV_grid.h5 Similar large patchy differences as above, plus also the Met Office reference has apparently missing data at -9999K that is not present in the ECMWF file.
	AMSR-2	Part 1: 6.4s Part 2: 3.0s BUFR files (generic) All values identical to reference except typicalSecond, typicalTime, minute and second, which always vary, e.g. 48 min 21 sec (met office) to 47 min 46 sec (ECMWF). Established that the set of reference results provided was not up to date. The time stamp was corrected during Met Office validation testing, to be consistent with EUMETSAT. *A new set of reference results will be generated before release.* BUFR files (grid) All values identical to reference except typicalSecond, typicalTime, minute and second, which always vary. Here, minute and seconds are vectors, except for apparently the first BUFR subset of the Met Office reference output, where minute is a scalar. All contain differences between reference and ECMWF. HDF file (generic) All contents identical to reference, apparently to machine precision, except for minute and second, where almost all values are different. HDF file (mapped) As for generic file, identical except minutes and seconds
	GMI	Part 1: 0.4s Part 2: 9s BUFR files S1 bufr file identical to reference (according to bufr_dump) except that spacecraftYaw value is 179.96 in Met Office reference and 79.96 in ECMWF Bug found: see section 5.4 S2 bufr identical except: Met Office ECMWF spacecraftRoll 0 0.1 spacecraftPitch 0 0.4 spacecraftYaw 179.96 79.96 Bug found: see section 5.4 H5 generic files S1 channel 1 longitude, latitude and brightness temperature visually identical (plotting H5 using IDL) and very similar to reference image generated using Cartopy H5 gridded Channels 1, 9, 10 longitude, latitude and brightness temperature compared to reference and were visually identical (and similar to Cartopy reference).

<table>
<tr><td></td><td></td><td>Numerically, differences of up to 1e-4 Kelvin were found.

Second, Minute, Hour, Day, Year, SatelliteZenith, SatelliteAzimuth all identical to reference

Satellite zenith of e.g. 52.90 is correct for theS1 channels but incorrect for S2 (where it is 48 degrees) A software change was implemented, see section 5.4

Satellite azimuth of -10000 is obviously incorrect (presumably missing) but of critical importance. It needs to be correctly extracted. Cannot be fixed because the satellite azimuth is missing from the hdf5 source files and also from EUMETSAT's BUFR product which we use as reference.</td></tr>
<tr><td></td><td>MWRI</td><td>Part 1: 3s
Part 2: 36s

The reference_products directory contains 2 additional files that are not generated by the test case:

FY3C_MWRIA_GBAL_L1_20180613_0715_010KM_MS.HDF_grid.h5.bufr
FY3C_MWRIA_GBAL_L1_20180613_0715_010KM_MS.HDF_grid.h5.h5
These were a left-over from earlier tests. A new set of reference results will be generated.

Plain BUFR file

ECMWF generated file is completely identical to the reference.

Gridded BUFR file

All apparently identical apart from rare differences in TB of around 0.01K. OK

HDF5 generic

Apparently identical to reference apart from numerical differences in TB of around 1e-4K

HDF5 gridded

As generic, all appears identical except numerical differences in TB of around 1e-4K</td></tr>
<tr><td>Timing testing</td><td>Compare with Met Office timings</td><td>Generally up to 2x faster than Met Office reference. Good.</td></tr>
<tr><td>Documentation testing</td><td>Review User Manual</td><td>Generally extremely good user manual, with a few minor issues:

- In the area of installation instructions. Here instructions make use of environment variables ECCODES_INSTALL_DIR and HDF5_INSTALL_DIR which are only defined if the instructions for locally building these libraries have been followed. Where installations already exist, and the local build has not been done, it is not clear at which level of the package installation tree these variables are supposed to point.
The documentation has been updated to make this clearer.

- When placing on grid, describe what happens to the longitude and latitude (see also additional comments) The code has been changed so that original lat/lon values are output, and the descriptions in the User Manual of mwipp_write_bufr_usergrid_mod.f90 and mwipp_write_hdf_usergrid_mod.f90 have been modified accordingly.</td></tr>
<tr><td></td><td>Review Release Note</td><td>About the right length and useful</td></tr>
</table>

	Review test case instructions	Minimal README.txt but sufficient. Consider listing the names of the expected output files This has been added to the README.txt. Consider automating some of the testing, e.g. using diff and bufr_dump, in cases where there should not be numerical differences compared to the reference (e.g. where averaging is not involved). This could be considered for a future release. Have added a comment that BT differences should be of order 0.01K.

Additional comments from the beta testers should be recorded below.

- Mapping onto an external grid is primarily useful for thinning data. Therefore it is preferable if the observations optionally keep their original longitude and latitude, rather than taking that of the reference grid. This enables more accurate simulation of the thinned observations. Agreed. It is better to keep the original lat/lon. Code has been modified (*mwipp_write_bufr_usergrid_mod.f90* and *mwipp_write_hdf_usergrid_mod.f90*)

- When thinning, a regular lat/lon grid over the poles is unhelpful as it will generate far too dense a coverage of polar data. Here the reduced Gaussian grid is perfect. Hence, consider supplying a set of standard reduced Gaussian grids to give users some easy options to define their thinning grids. It is not straightforward for users to define these grids, or to know how to extract them from the ECMWF archives.
  Agreed. Propose to make available a set of grids on the NWPSAF ftp site, alongside the test cases.
- The mapping onto a defined grid will certainly be useful, but it does not replicate all the requirements of the ECMWF all-sky processing, where observations are first mapped onto points of a user-defined grid and then all obs associated with a grid point are averaged to create superobs, possibly using a maximum distance from the grid-point, and with many other refinements (e.g. time bins, ability to thin multiple sensors together). Of all these, the most important difference is that the gridding cannot be used to average, since this has to be done beforehand using weighted spatial averaging on the full input data. Something broadly comparable to the ECMWF approach could be achieved with MWIPP if unweighted spatial averaging were available, but using a distance cutoff. An unweighted averaging has been added, see section 5.4
- Mapping onto a grid is one easy way of combing the GMI S1 and S2 swaths but care is needed with regards the zenith angle, which is ~53 degrees for S1 but ~48 degrees for S2. Ideally the zenith angle needs to become a function of channel. Current MWIPP output for GMI is misleading, and there will be large errors in simulating the S2 channels if the wrong zenith angle is assumed. Note that RTTOV cannot handle multiple zenith angles per location, so at ECMWF we have to call it in a loop over S1 and S2 channels, using a different zenith angle each time. See section 5.4

5.3 Beta test report from UW SSEC

Table 5: Test setup: SSEC

Institute	UW SSEC
MWIPP version	0.1
ecCodes version	2.8.2
hdf5 version	hdf5-1.10.1.tar.gz
Machine name	milk (but in a Debian-based docker container)
uname -a	Linux 405cb2579ff2 4.9.93-linuxkit-aufs #1 SMP

	Wed Jun 6 16:55:56 UTC 2018 x86_64 GNU/Linux
Memory cat /proc/meminfo	MemTotal: 263729316 kB MemFree: 22890352 kB MemAvailable: 250573948 kB
Processors lscpu	CPU(s): 48 Thread(s) per core: 2 Core(s) per socket: 12 Socket(s): 2

Table 6: Beta test results: SSEC

Category	**Specific test**	**Beta tester comments**
Coding inspection	Review code, at the discretion of the beta tester	none
Compilation testing	Tester to select compiler	GNU Fortran (GCC) 8.2.0
User test cases	SSMIS	OK. real 0m45.212s user 0m50.222s sys 0m3.502s
	AMSR-2	OK. real 0m11.633s user 0m8.325s sys 0m0.814s
	GMI	OK. real 0m9.096s user 0m7.554s sys 0m1.504s
	MWRI	OK. real 0m34.852s user 0m38.376s sys 0m0.853s
Timing testing	Compare with Met Office timings	Reasonable? Yes, they are faster.
Documentation testing	Review User Manual	Clear and thorough. Some notes below.
	Review Release Note	Clear & concise. Noticed no issues.
	Review test case instructions	Fairly clear. As someone who's not very familiar with these instruments, I would have liked some instructions for how to compare against the reference output (what is a reasonable epsilon for brightness temp? machine precision? .01 K? etc.), or for the script to do this and return pass/fail or print a report to send back.

Additional comments from the beta testers should be recorded below.

- The ecCodes & hdf5 script use: FC, ECCODES_INSTALL_DIR, HDF5_INSTALL_DIR but the MWIPP configure script uses: FC, ECCODES_LIB, HDF5_LIB. It'd be nice for it to be consistent or, otherwise, for the configure script to look for the INSTALL_DIR flags and use that to generate the LIB locations.

- The hdf5 build script has –enable-fortran2003, this was crashing the build and doesn't seem to be a supported option for the hdf5 version that's downloaded (1.10.1) so I removed it and proceeded on. User Guide has been updated to say that --enable-fortran2003 is only necessary for v1.8.x.

- I also had to remove the check in MWIPP's configure script lines 115 to 126, though, or it wouldn't pass. These lines have been removed.

- The MWIPP configure script requires ksh to be installed on the system. I don't think this is documented, and it wasn't installed by default on my machine. It runs OK under ksh or bash. I have changed the script to default to bash and have mentioned this in the documentation.

- MWIPP's Makefile.ARCH appears to have "MAKE = gmake" hardcoded; I had to replace this with regular "make" for my system. The build script has been updated to accept either gmake or make. It tests to see which one is available on the system.

- When trying to run the binaries without the included scripts, I noticed that they don't have -h or --help flags. These would be useful, though the user guide does do a good job of explaining the required options. Actually the help info is printed if no argument is supplied. I have changed them so that the help info is also printed if only 1 argument is supplied (all the legitimate options require 2 arguments).

- I noticed most binaries segfault if provided input that doesn't exist (in my case, a link to a non-existent file). I have added tests for the existence of the input file.

- The quicklooks script MWIPP/bin/mwipp_quicklook.py uses /bin/env to find the python interpreter. This didn't exist on my system. I think /usr/bin/env is more common. This has been changed.

- Section 3.2's in the User Guide references mwipp_main.exe, I think it should be ssmis-main.exe. Confirmed. The User Guide has been corrected.

5.4 Issues arising from beta testing

Code changes are needed for the following issues:

- Satellite zenith angle is different for the GMI S1 and S2 channels, and this needs to be indicated in the products that are mapped to the user grid (BUFR and hdf5). This is fixed in r1820, as follows:
    - The satzen and sataz arrays in the mapped h5 product can be multi-dimensional. The dimensions are set in the top-level routine (gmi_main.f90). An attribute indicates the start and end channels. No change for other instruments.
    - For the BUFR product, an additional replication has been added, so that we can have more than one set of angles. The start and end channels are given via additional descriptors. See the User Guide for details.

- Incorrect satellite yaw for GMI. This was a bug. Fixed in r1819.
- Keep original lat/lon in mapped products. Implemented in r1824.
- Unweighted averaging option, requested by ECMWF. Implemented in r1821. User Guide updated.
- In Makefile.ARCH, use "make" if "gmake" is not available. (SSEC suggestion). Implemented in r1822.
- Add error traps for non-existent input files, and update script interpreter lines. (SSEC). Implemented in r1823.
- Noted by developer: SSMIS output products, except for UPP BUFR, do not contain satellite zenith angle. This could be created from the "base point" data, but will require significant effort. Propose to postpone this to an update release.

5.5 Regression test

A regression test is necessary in order to ensure that the reference results in the user test cases are consistent with the final version of the software.

Table 7: Test setup (regression)

MWIPP version	1828 2018-11-12 14:11:46
ecCodes version	2.9.0
hdf5 version	1.8.16
Machine name	eld030
`uname -a`	Linux eld030 2.6.32-754.3.5.el6.x86_64 #1 SMP Thu Aug 9 11:56:22 EDT 2018 x86_64 x86_64 x86_64 GNU/Linux

Memory
`cat /proc/meminfo`
	`Total: 6116832 kB`
	`MemFree: 590124 kB`
	`Inactive: 1351924 kB`

Processors
`lscpu`
	`CPU(s): 8`
	`Thread(s) per core: 2`
	`Core(s) per socket: 4`
	`Socket(s): 1`

Table 8: Test results (regression)

Category	Specific test	Performed by	Date	Result
Re-run test cases	SSMIS, gfortran	NCA	12.11.18	Ran OK. Updated the 00README file as recommended by beta testers. Outputs saved as reference.
	AMSR-2, gfortran	NCA	12.11.18	Ran OK. Updated the 00README file as recommended by beta testers. Outputs saved as reference.
	GMI, gfortran	NCA	12.11.18	Ran OK. Updated the 00README file as recommended by beta testers. Outputs saved as reference.
	MWRI, gfortran	NCA	12.11.18	Ran OK. Updated the 00README file as recommended by beta testers. Outputs saved as reference.

6. VALIDATION TEST RESULTS

This section documents the validation tests run at the Met Office under the configuration shown in section 5.1.

6.1 SSMIS

The SSMIS test case was run in a number of different configurations and the results compared with SSMIS_PP.

The run time of the SSMIS_PP Averaging module was 6 seconds (gfortran compiler)

Reference	Configuration	Run-time	Comment / result
1	nweights = 100, min_weight = 0.01, all channels	20s	Default configuration. See Figure 1
2	nweights = 200, min_weight = 0.002, all channels	24s	Same configuration as SSMIS_PP. See Figure 2
3	As 1, but for two channels only	20s	Confirmed that the correct channels were processed
4	always_average = .true.		Confirmed that averaging was done in cloudy regions.
5	rain_averaging_threshold = 1.0		Confirmed that BTs in cloudy regions were a close match to SSMIS_PP – but this is not a good configuration scientifically

Maps of the differences with respect to SSMIS_PP are shown in Figure 1 and Figure 2.

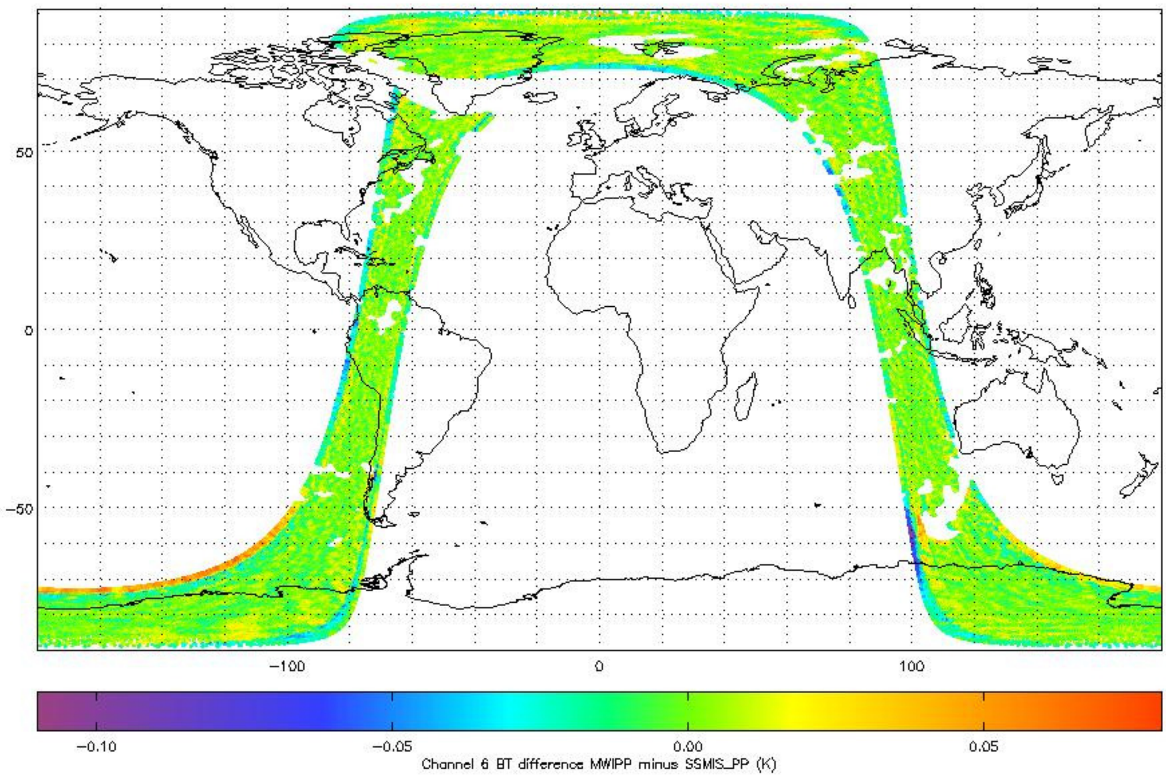

Figure 1: Channel 6 BT difference from SSMIS_PP for configuration 1, for non-raining areas

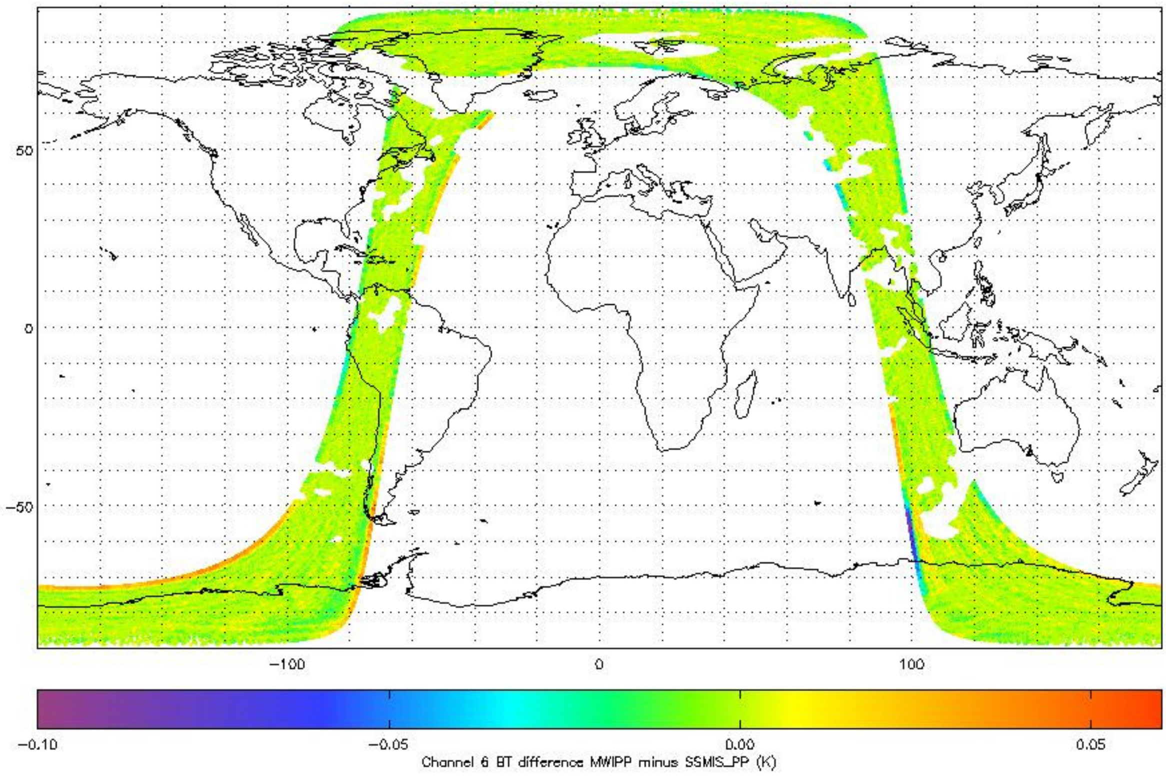

Figure 2: Channel 6 difference from SSMIS_PP for configuration 2, for non-raining areas

Channel 6 (57.29 GHz) was chosen for these comparisons because it is the lowest sounding channel that is assimilated operationally at the Met Office (channels 1-4 are faulty on SSMIS F17). Figure 1 shows that there are some small systematic differences between MWIPP and SSMIS_PP at the swath edges, which are still present when the number of points in the averaging is increased to 200 (Figure 2). The most likely explanation for these differences is that the weighting coefficients for SSMIS_PP were determined many years ago using SSMIS F16, whereas the weighting coefficients for MWIPP are determined using real data, in this case for F17. These weights for spot 1 are plotted in Figure 3. There are indeed some subtle differences, e.g. in SSMIS_PP the spot offset extends to 9 spots whereas MWIPP it only reaches 8. It should be noted that averaging at the edge of swath is inevitably slightly inaccurate because the centre of gravity of the points to be averaged does not coincide with the spot under test.

Are these differences significant? In Figure 4 we show that differences with respect to *unaveraged* data are more than an order of magnitude larger than the differences between MWIPP and SSMIS_PP. Looking closely at the data (interactively) there was no evidence that the SSMIS_PP was any better or worse than MWIPP. We conclude that the differences are acceptable.

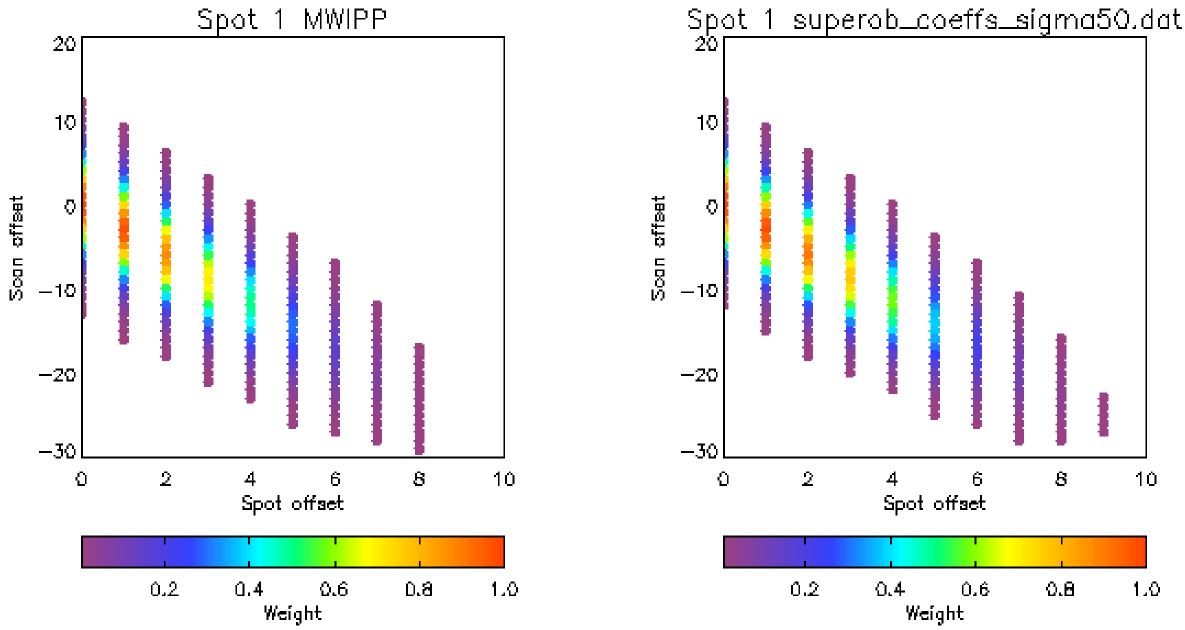

Figure 3: Left: weights for spot 1 from MWIPP, right: weights for SSMIS_PP

Note that the rain-flagged areas are slightly larger in Figure 2 than Figure 1. This is to be expected, as the search area is wider.

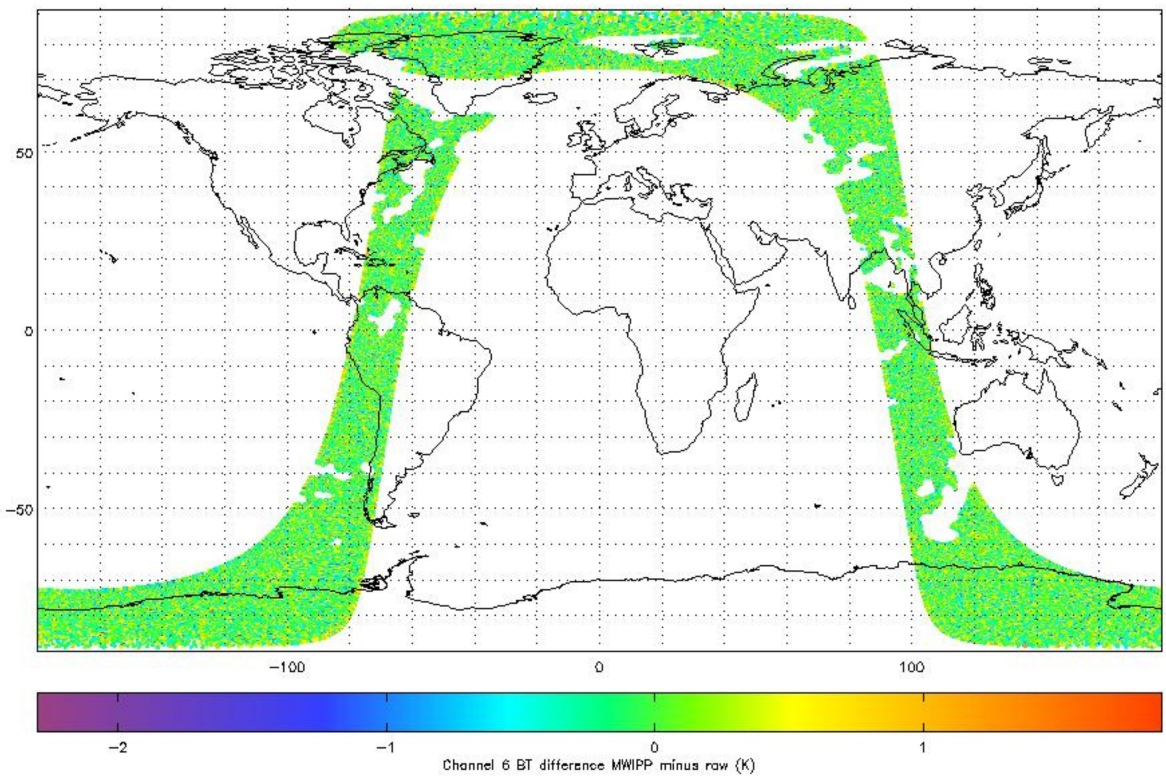

Figure 4: BT difference compared with raw, unaveraged data. Note the differences are more than an order of magnitude greater than the differences between MWIPP and SSMIS_PP.

Similarly the extent of "coast" is larger in configuration 2 (Figure 6) than configuration 1 (Figure 5 and zoomed-in version at Figure 7). The extent of the coast in configuration 1 is about 190km. This seems quite generous bearing in mind that the amplitude of the Gaussian averaging has fallen below 0.001 at this distance and is 0.01 at a distance 150km.

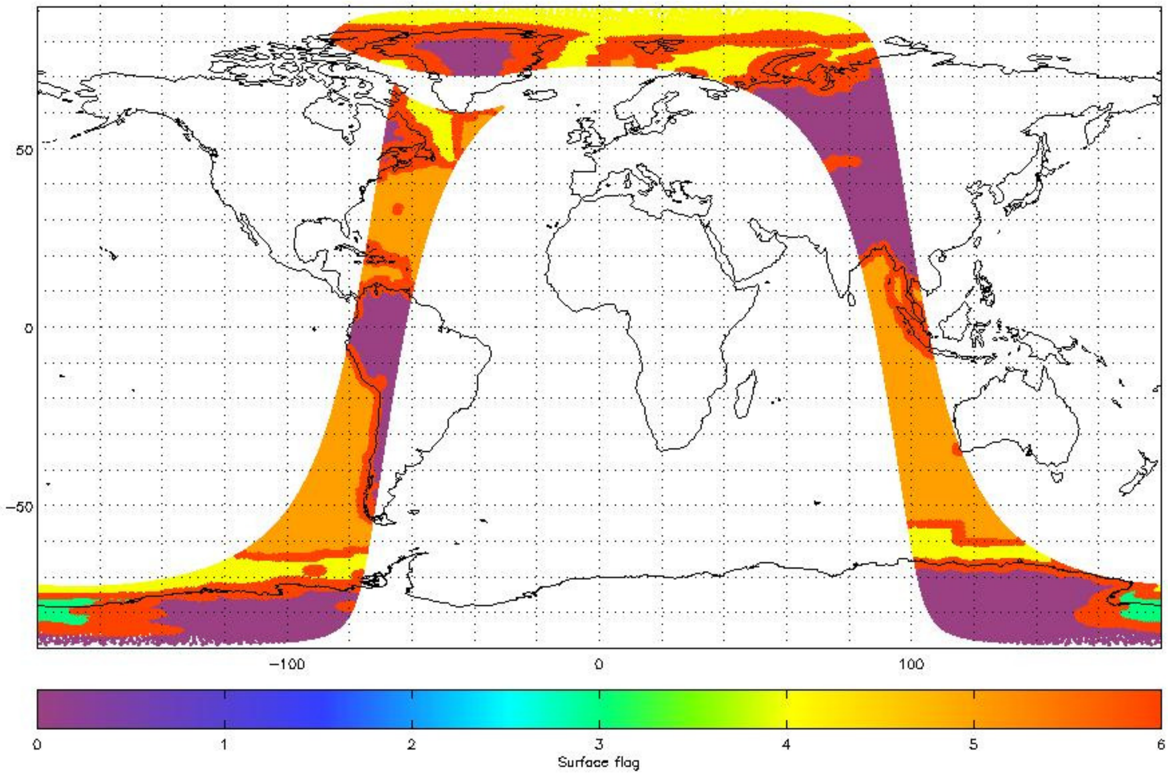

Figure 5: Surface flag for configuration 1 (100 samples averaged)

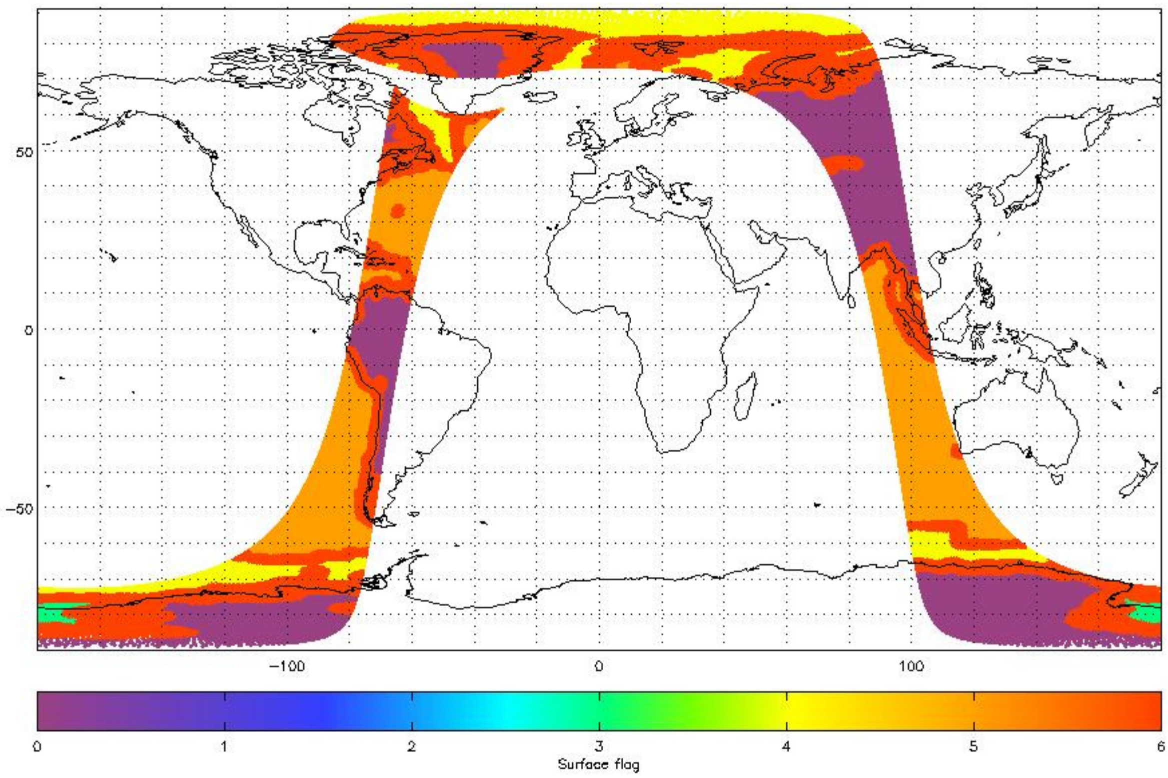

Figure 6: Surface flag for configuration 2 (200 samples averaged)

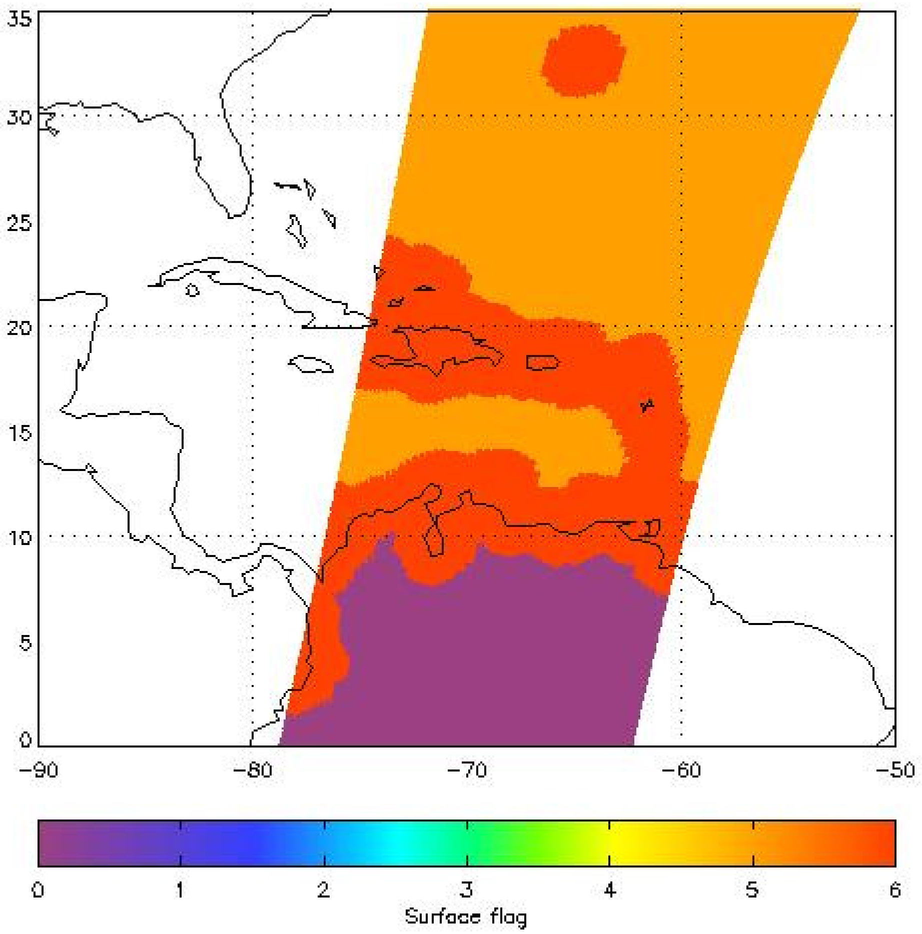

Figure 7: Zoomed in view of Figure 5 (100 samples averaged). The extent of the coast is about 190km.

Looking at the run times for MWIPP: in configuration 2, the run time is a factor 4 longer than SSMIS_PP (slightly shorter in configuration 1). The main reasons identified were:

- ecCodes is significantly slower than the Met Office BUFR library that is included in SSMIS_PP. However, there are good reasons for choosing ecCodes for this deliverable:
    - ➢ long-term support
    - ➢ easy to use
    - ➢ widely used internationally
    - ➢ licensing arrangements
- The loop in modify_btemps subroutine takes about 8 seconds with 200 points and 4 seconds with 100 points. It is slightly more complex than the equivalent in SSMIS_PP, which has fewer processing options and is not a generic routine. But it operates on the same principle.

24 seconds to process a full orbit (100 minutes) of data is not unreasonable. Request a waiver on the run time.

6.2 AMSR-2

AMSR-2 hdf5 data from JAXA were processed to BUFR and compared with EUMETSAT data. Input files:

- GW1AM2_201809240134_189A_L1SGBTBR_2220220.h5
- W_XX-EUMETSAT-Darmstadt,SOUNDING+SATELLITE,GCOMW1+AMSR-2_C_EUMP_201809240140_33791_33792_110_L1B.bin

The JAXA data are in the form of half-orbits, the EUMETSAT data are 5-minute granules.

Comparison is awkward because there are 4 scans per message in each case, and the message boundaries do not in general coincide. However, in the above case they do coincide.

The first message in the EUMETSAT file was compared with the 61st message in the JAXA file, analysed using *bufr_dump*. Differences are recorded in Table 9. Values quoted are for the first spot in the message, unless stated otherwise.

Table 9: Comparison between MWIPP and EUMETSAT BUFR files for AMSR-2

Variable	MWIPP	EUMETSAT	Comment
typicalSecond	7	0	All the "typical" values are the same in the EUMETSAT file; in MWIPP they change with each message. Both approaches are valid.
typicalTime	014007	014000	As above
second	7	6	Checked with the original hdf file. MWIPP is rounding to the nearest second, EUM is probably rounding down. Not an issue.
scanlineNumber	241	1981	MWIPP starts at 1 for the first scan in the file, which is reasonable. Not present in the hdf.
solarZenithAngle	60.99	Missing	There is a problem with the solar angles in the source file. They do not agree with the NOAA Solar Calculator[1]. This is unlikely to be an issue for users as solar angles are not generally used for microwave instruments. MWIPP approach is simply to report the supplied values.
solarAzimuth	100.75	280.75	Difference of 180°, however neither is correct, as noted above.
satelliteZenithAngle	55.18	235.18	The MWIPP value is correct (consistent with WMO Oscar database). EUM value differs by 180° and is unrealistic.
bearingOrAzimuth	205.09	25.09	The issue here appears to be whether to report the satellite azimuth as seen from earth (MWIPP) or the earth azimuth as seen from satellite (EUM). Not clear in the BUFR table definition. The value reported by MWIPP is same as the value in the source file.

In the hdf5 source file, an inconsistency was noted between the header attribute "ObservationStartDateTime" and the dataset "Scan Time" (in seconds). The following comment

[1] https://www.esrl.noaa.gov/gmd/grad/solcalc/

was inserted into the source code to document this. It is not an issue for MWIPP output products, therefore not necessary to mention it in the user documentation.

```
! The ObservationStartDateTime is given as "Z" (i.e. UTC), however it is 36
! seconds ahead of the first scantime. If it was really UTC, and the scantime
! was in TAI then the header would be 36 seconds behind. Therefore we ignore the
! header time as being unreliable and assume the scantime is UTC. This is what
! EUMETSAT have done in their encoder. The result appears to be consistent with
! the positions given at http://www.ssec.wisc.edu/datacenter/orbit_tracks.html
```

6.3 GMI

GMI hdf5 data from JAXA were processed to BUFR and compared with EUMETSAT data. Input files:

- 1C.GPM.GMI.XCAL2016-C.20180410-S122640-E123138.V05A.RT-H5
- W_XX-EUMETSAT-Darmstadt,SOUNDING+SATELLITE,GPM+GMI+S1_C_EUMP_201804101226_122640_123138_L1C.bin
- W_XX-EUMETSAT-Darmstadt,SOUNDING+SATELLITE,GPM+GMI+S2_C_EUMP_201804101226_122640_123138_L1C.bin

The data are in the form of 5-minute granules. Each file was analysed using *bufr_dump*. Differences are recorded in Table 10. None of the differences are significant.

Table 10: Comparison between MWIPP and EUMETSAT BUFR files for AMSR-2

Variable	MWIPP	EUMETSAT	Comment
typicalSecond	40	0	All the “typical” values are the same in the EUMETSAT file; in MWIPP they change with each message. Both approaches are valid.
typicalTime	122640	122600	As above
sunGlintAngle	95, …	0	All sun glint angles are missing in EUMETSAT data. Copied unchanged from the source for MWIPP.

7. SUMMARY

The status of the requirements traceability matrix at the conclusion of testing is shown in Table 11.

Table 11: Requirements matrix status at the conclusion of testing

Identifier	Requirement in Product Spec	Testing method	Status
MWIPP1	Documentation is clear, understandable and complete	Beta testing	Pass. The comments of the beta testers have been addressed.
MWIPP2	Code conforms to the requirements of NWPSAF-MO-SW-002: commented, understandable and modular	Inspection	Pass
MWIPP3	Any necessary external libraries are freely available	Inspection (e.g. examine the	Pass

		corresponding web sites for the external libraries)	
MWIPP4	Code builds with no errors on a 64-bit Linux PC, running a current operating system such as CentOS7 or RHEL7. More than one Fortran compiler shall be tested.	Test	Pass. Tested gfortran and ifort.
MWIPP5	Ingest BUFR files for SSMIS and AMSR-2	Test	Pass. Included in the SSMIS test case
MWIPP6	Ingest native-format files for AMSR-2	Test	Pass. Included in the AMSR-2 test case.
MWIPP7	Spatial averaging capability (replicates SSMIS_PP)	Test	Pass. Included in the SSMIS test case
MWIPP8	Map to user-defined grid	Test	Pass. Included in all 4 test cases.
MWIPP9	Creation of BUFR output files for SSMIS and AMSR-2 (optionally GMI and MWRI)	Test	Pass. Included in all 4 test cases.
MWIPP10	Creation of hdf5 or NetCDF4 output files	Test	Pass. Hdf5 file generation is included in all 4 test cases. NetCDF4 is not implemented in this release since none of the instruments provide NetCDF4 input datasets.
MWIPP11	Run times are documented in the test log, and are comparable with SSMIS_PP when processing SSMIS.	Test and inspection	Run times are slower than SSMIS_PP by a factor ~4, largely due to the use of ecCodes rather than Met Office BUFR library. But considered acceptable for operational use. Waiver requested.
MWIPP12	Test cases for the users exist, have clear instructions and run correctly	Beta testing	Pass. The test case instructions have been updated following beta testers feedback.

Other points:

- For a future release, consider extracting satellite zenith angles for SSMIS, as this quantity is missing for re-mapped products.
- For a future release of the test cases, consider automation of the comparison with reference results.
- Solar angles are incorrect for AMSR-2 data. But this is a fault in the source data, not MWIPP.