
IFS DOCUMENTATION – Cy50r1
Operational implementation 12 May 2026

PART I: OBSERVATIONS

© Copyright 2020-

European Centre for Medium-Range Weather Forecasts
Shinfield Park, Reading, RG2 9AX, UK

The content of this document is available for use under a Creative Commons Attribution 4.0 International Public License. See the terms at <https://creativecommons.org/licenses/by/4.0>.

REVISION HISTORY

Changes at CY43R1

- Major rewrite and restructure of the observation documentation reflecting OOPS developments for CY43R1, removal of outdated material and identification of deprecated areas of code.

Chapter 1

Overview: the observation world

Table of contents

1.1	Introduction
1.2	Concepts
1.2.1	Observation groupings, reports and datums
1.2.2	Observation sets
1.2.3	Timeslots
1.2.4	Events and status
1.2.5	Screening trajectories
1.2.6	Observation Database usage in IFS
1.2.7	Parallel aspects (scalability)
1.3	Dataflow through processing and screening
1.4	Observation groupings in the IFS
1.4.1	MARS reporttype and group
1.4.2	IFS Obstype and codetype
1.4.3	SQNO - [DEPRECATED]
1.4.4	IFS Physical variable: varno
1.4.5	Observation operator codes: NVAR
1.4.6	Satgrp_table and friends
1.4.7	'Area' type - [DEPRECATED]
1.4.8	VarBC bias groups

1.1 INTRODUCTION

The handling of observations is one of the most complex parts of the IFS data assimilation system. Observations come in diverse forms and, depending on their type, they are treated in very different ways. The observations themselves (y) and the observation operator $H()$, which converts from the model state x to an observation equivalent $H(x)$, are just part of a wider infrastructure of processing, quality control, thinning, and assigning observation errors. The metadata associated with each observation can run into hundreds of variables, all needing to be stored in an observation database (ODB). Thus the observation world encompasses:

- **Observation operators** as described in [Chapter 3](#). The input to the observation operator is the ODB, plus the the model state at observation locations. Observation operators, as well as generating $H(x)$, may put additional information into the ODB in some cases including observation error and quality control decisions.
- **Model state interpolated to observation location** or **GOM.PLUS** is described in Part 2 (Data Assimilation). This contains vertical profiles of atmospheric variables (temperature, humidity etc.), surface variables and any other information required from the model. Using the '2D GOM' facility, instead of a single vertical profile, the GOM.PLUS can contain a slant-path, a set of profiles along a limb path (e.g. for GPSRO) or just all the nearby model columns (which is used in the Bayesian radar retrieval at Météo-France).
- **ODB** is a collection of data formats and metadata describing the observations.

- In the core of the IFS, processing is based on ODB-1, a hierarchical table format capable of running in a parallel environment, manipulated and accessed using Structured Query Language (SQL). Documentation on ODB-1 can be found at <https://confluence.ecmwf.int/display/ODB/ODB+User+Guide>.
 - However, most interactions with ODB in the IFS use an additional layer of code known as ifsobs (see Sec. 1.2.6).
 - For some pre-processing, and for archiving to MARS, the ODB-2 format is used: <https://confluence.ecmwf.int/display/ODBAPI>.
 - The variables and codes used in the ODB are governed by a committee, and can be browsed at <http://apps.ecmwf.int/odbgov>.
 - At a more detailed level, some parts of the ODB contents are described in the current documentation, mainly in [Chapter 6](#).
- **Observation processing** is a chain that leads from the ingestion of raw data (as BUFR files and other data formats) to the archiving of ODBs in MARS containing diagnostics from the assimilation run (traditionally known as feedback files). Apart from running the observation operators in the data assimilation, observation processing has two main components that can be described as ‘preparation’ and ‘screening’, described next.
 - **Preparation** generates all the parameters needed to run the observation operator in the assimilation system. This is everything from format conversions (e.g. BUFR to ODB-1) to the retrieval of surface emissivities or cloud top pressure from satellite radiances. It is also necessary to apply coordinate transforms (for example from windspeed and direction to horizontal wind components) or to run a retrieval prior to assimilation (for example from backscatter to ocean surface wind). This stage may include simple thinning (such as discarding some satellite observations) This stage also sets up observation errors and other parameters such as the settings for quality control. All this is described in [Chapter 2](#).
 - **Screening** includes all the decisions whether or not to use a certain data. Screening decisions are made throughout the processing chain. They includes blocklisting and other quality control decisions that can only be made when a first guess is available, or in the presence of all observations (‘dependent’ screening). See [Chapter 4](#).
 - **Observation errors:** the diagonal part of the observation error matrix \mathbf{R} is created during observation processing. For conventional observations, it is done in **SUOBSERR**; for other observations it may be created in **DEFRUN** or the first time the observation operator is called, particularly if the observation error is situation-dependent. This observation error is eventually stored in the ODB for later use by the data assimilation system. Correlated parts of the observation error model are described in Part 2 (Data assimilation).
 - **Bias correction** Some bias correction is performed as part of pre-processing (such as for radiosondes, done in **TEMPIN**) but most is now handled inside the 4D-Var minimisation using variational bias correction (VarBC). This is more properly part of the data assimilation, so it is also left to Part 2.
 - **Other auxiliary control variables or TOVSCV** Similar to VarBC, other auxiliary control variables are used as sink variables for satellite radiance assimilation and/or observations for further assimilation in the ocean. The latter include ocean skin temperature and sea ice concentration, which are assimilated in the ocean data assimilation. Observation space auxiliary variables (TOVSCV) are partly described in Part 2, but their main documentation is here, in [Chapter 3](#).
 - **Variational quality control** This is part of the data assimilation algorithm, so again it is found in Part 2. However, it is the responsibility of the observation processing to provide the parameters that control these tests, such as for example the Huber norm parameters; these are stored in the ODB for use during the assimilation.
 - **Diagnostic Jo table** Although this is another Part 2 concept, the diagnostic Jo table has a detailed structure that breaks down observations into smaller groups, such as observation type and variable for the conventional observations (e.g. radiosonde wind) and satellite instrument for radiances (e.g. AMSU-A). This fine structuring of observations is known only by the observation world, and is described here.

As with the rest of the IFS, the observation world has built up over more than 40 years and is still in constant development. Processing takes place in different layers of code, from ancient parts of the original IFS, through to more recent OOPS encapsulations. We are constantly working to remove and refactor the older layers of code, but this is an ongoing task. Some areas and concepts are now 'deprecated', in other words they still provide important functions but they are expected to disappear in future cycles and they should not be involved in new developments. More broadly, the exercise of documenting code is a good way to understanding its inadequacies and indicating ways its design could be improved; thus it also tries to identify areas that could be rationalised in future cycles. To understand the code as it stands, it is worth knowing the history of such developments. For example, before the ODB was introduced, observations were stored in a data structure called CMA, and concepts from the CMA live on in many places, some good, others outdated and in the process of being refactored. The ability to call observation operators has now been encapsulated for OOPS, but areas such as setup and screening have not seen so much development.

1.2 CONCEPTS

1.2.1 Observation groupings, reports and datums

Observations are organised into a hierarchical structure which is reflected in ODB-1. At the broadest level, observations can be organised into meaningful groupings, such as the data from a sensor on one satellite, or all drifting buoys. There are a number of schemes for grouping used in the IFS, and many of these groupings overlap, since they make sense in some areas but not others. Lower down is the concept of a 'report', encompassing for example a single radiosonde ascent or a satellite observation that spans multiple frequencies. Each report thus contains multiple observations, often referred to as datums. Strictly, the plural of datum is data, but 'datum' or 'datums' emphasises that this is the lowest quantum of observational information. In the ODB each datum is associated with a vertical position (`vertco_reference.1@body`, given in terms of pressure, height or channel number) and a variable number (`varno@body`). Datums are grouped into reports to take advantage of their common characteristics: for satellite observations, we often loosely think of the report level as the 'observation'; the datums in the report (measurements at different frequencies) relate to a single location and were made using the same observational parameters, such as scan angle or instrument temperature. Hence it is more precise to talk about datums and reports than 'observations'.

1.2.2 Observation sets

To allow efficient parallel computation in the IFS, reports are usually processed together in a small groups known as sets. The maximum number of reports in a set is given by `NMXLEN` (`yomdim0`; default 511 but currently set to 128 under control of `ExpERT`, the experiment submission system). To balance the parallel processing, sets are handled in order from the most computationally expensive to the least. The observation sets may span several 4D-Var timeslots, though in the current operational context this is not the case. Sets are generated by the class `OBSOP_SETS` with their detailed organisation being handled by `ECSET`. Satellite observation sets (e.g. radiances, AMVs, limb radiances, scatterometer) must not contain data from more than one instrument. This is controlled by the misnamed 'area' parameter, which for radiance data is an indicator of satellite ID and instrument, determined in `SUOBAREA`.

1.2.3 Timeslots

Timeslots are sub-divisions of the 4D-Var window into which observations are grouped for processing. For interpolation from model to observation space, one timestep of the model is chosen to represent the model state in that timeslot. Timeslots are useful because they isolate the observation grouping from the model time resolution, which varies within 4D-Var. Currently with a length 15 minutes, timeslots could be aligned in future to simply use the most appropriate model timestep for each observation.

1.2.4 Events and status

The results of observational processing are stored in the ODB, at the report and datum level, in columns called `REPORT_STATUS`, `DATUM_STATUS`, `REPORT_EVENT1` and `DATUM_EVENT1`

(Secs. 6.7 and 6.8). The status summarises whether (and how) the observations should be used in the assimilation system; the event columns are bitfields that record the causes of data rejection, plus some other processing events. Many processing decisions are unique to a particular observation type. For this reason, the `REPORT_EVENT2` and `DATUM_EVENT2` columns can be customised according to the observation type, although in practice this facility is rarely used. All-sky radiance observations use their own event bitfields called `DATUM_TBINFO`, `DATUM_TBFLAG` and `DATUM_TBERROR` (Sec. 3.3.3).

1.2.5 Screening trajectories

Most observations rely on a ‘screening trajectory’ to perform one-off initialisations that may require the model’s first guess, and there is also the pre-processing chain which contributes to these initialisations and data selections. Within the main IFS, all these screening decisions are activated under the flag `LSCREEN = .TRUE`.

The concept of continuous data assimilation has changed the way in which this is done. Traditionally, screening was something done once, based on the background trajectory, also known as the first trajectory. `LSCREEN` was true in the first trajectory and nowhere else. But now with either the `LCONTINUOUS_DA` or `LREScreen_*` switches, all the preprocessing and screening may be done every trajectory of the incremental data assimilation, except for the final analysis trajectory (where `LSCREEN` is false).

Given `LSCREEN` triggers such a wide range of behaviour from setting observation error models, running prior retrievals, as well as true screening, there can be confusion over what ‘screening’ means. This documentation tries to make the distinction between ‘preparation’ (some of which takes place when the screening is being done, but is not screening) and true screening, which selects data to go into the assimilation algorithm, and is applied at or after the stage the observation operator is run.

1.2.6 Observation Database usage in IFS

ODB-1 is an MPI-parallel distributed in-memory, hierarchical database developed at ECMWF. The IFS uses ODB-1 to access observational data and store feedback data (for example, to store first guess departures and information on whether a particular observation was rejected by screening).

Typically, when IFS needs to access observational data, an SQL query is run to retrieve a subset of the data. The queries themselves can be found in the `odb/dll` directory. The results from the query are returned in 2D arrays. The IFS then uses and manipulates these data arrays, before the modified data is put back into the ODB again. Often, separate queries are run for the report-level data (stored in `ROBHDR`) and the datum-level data (stored in `ROBODY`). In a typical screening trajectory, each MPI task performs $O(10,000)$ separate SQL queries on the database.

Although IFS uses ODB-1, there are no direct calls to the ODB library from the IFS source code. Instead, IFS accesses ODB through an interface layer called `ifsobs` which was first introduced in `CY43R1` and gradually rolled out across IFS in cycles up to `CY48R1`. The interface provides an object-oriented handling of the observation databases. In principle, `ifsobs` allows different backends to be plugged in, of which ODB-1 is just one option, while keeping the IFS code unchanged. SQLs are run using the `SELECT` method of the database object and data is returned in a `DBASE_VIEW` object - essentially a data frame containing the data alongside metadata on the column names and types. Modified data can then be saved into the database by passing the modified views back into the database object via the `PUT` method. The `DBASE_VIEW` objects are generic and can store any data, regardless of the schema. However, for convenience inside IFS an `IFS_DBASE_VIEW` object (defined in `arpifs/module/ifs_dbase_view_mod.F90`) wraps the abstract `ifsobs DBASE_VIEW` class and adds pointers with names matching the ODB column names which point to the relevant data in the data array. This allows users to interact with ODB column data using the syntax `<VIEW_NAME>%<TABLE_NAME>%<COLUMN_NAME>(IROW)` e.g. `ROBODY%BODY%OBSVALUE(I)`. Further documentation of `ifsobs` is provided in the `ifs-source` repository itself linked from `ifsobs/README.md`.

1.2.7 Parallel aspects (scalability)

The IFS is multi-process multi-threaded parallel. For running the forecast model, each process holds the state variables for a small part of the globe. In contrast observations are (roughly randomly) distributed across all processors with the aim of avoiding any geographical link. Each processor works with a selection of observations from across the globe. This is why the interpolation from model space to observation space (the GOMs) is one of the more demanding parallel-processing tasks in the IFS. An alternative strategy of keeping the relevant observations on the processor that owns the model data would fall down in two ways: First, a horizontal interpolation, and even more so a limb or slant path interpolation, will often need access to model data located on multiple processors, so multi-process communication would be needed anyway. Second, the computational burden would be unequally distributed, concentrated in whichever processors have satellites overhead on that particular timestep: maybe a dozen out of hundreds.

Under the IFS data assimilation code, which is still functioning and still used operationally for FSOI computation (forecast sensitivity to observation impact) the top level for running the observation operator is **TASKOB** (or its tangent-linear and adjoint equivalents). This loops over the available observation sets on that processor. This loop is multi-threaded, which means that many sets will be processed in parallel. Moreover, given that the IFS is also multi-process parallel, each processor will be handling different sets of observations. Running the observation operator is thus almost perfectly scalable, as long as the workload is evenly distributed across all processors and threads.

Under IFS, the next level down is **TASKOB.THREAD** which handles the parallel parts of the observation loop. For each observation set, it calls **YDGOM5%MODEL_AT_OBS** to get the GOM_PLUS (the model state at observation locations) and then calls the per-set observation operator **HOP**, which is where [Chapter 3](#) (Observation operators) takes over. The screening uses a similar framework of a parallel loop over observation sets.

Under the OOPS data assimilation code, which is default for all other applications than FSOI, the top level of the observation operator is **ALLOBS.OPER.MOD** with **OBSEQUIV** and **OBSEQUIV_MT** providing the same functions and basically doing the same thing as **TASKOB** and **TASKOB.THREAD**, ultimately calling the universal observation operator **HOP** in multi-thread multi-task parallel. Similarly, tangent-linear and adjoint versions of these routines also exist to support the computation of updated cost functions and gradients as part of the minimisation.

1.3 DATAFLOW THROUGH PROCESSING AND SCREENING

Observation processing happens at many different stages during the journey through the IFS. For computational reasons it can be helpful to reduce observation numbers early on, when the data is stored compactly, and before too much computational time has been expended on unwanted data. However, the cost of observation processing in the current operational system is small, and these computational limitations are not as important as they have been ([Lean et al., 2016](#)). Also, much of the screening needs to be left until later stages because it cannot be done in isolation. It is often necessary to have access to:

- The first guess (FG) model state
- The first guess departure, i.e. $\mathbf{y} - H(\mathbf{x})$
- The results of screening decisions from other observations

The last of these points leads to an important distinction in the IFS: between ‘independent’ and ‘dependent’ screening decisions. The dependent decisions are things like checking for redundant observations and thinning, where a selection is made among all the observations that passed earlier quality checks. For example, clear-sky satellite radiance thinning takes as input all observations that passed quality control and blocklisting, and it selects observations with the warmest brightness temperature, as a way of filtering otherwise undetected cloud.

There is thus an opposition between wanting to discard data early for performance reasons, and keeping it until enough supporting information is available, such as the FG model state, or the FG departure. For this reason, and also to perform actions in the easiest place, as well as just for historical reasons, there is

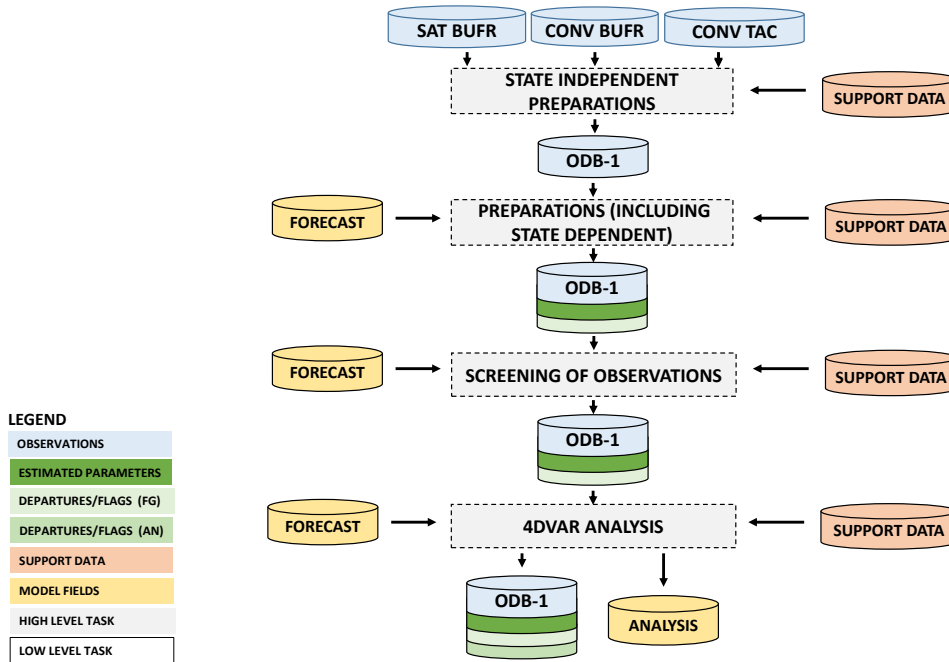


Figure 1.1 Simplified IFS observation processing flow diagram, with the disc icons representing data stores and the rectangles representing processing stages. Further details can be found in Fig. 2.1 for the state-independent preparations, Fig. 2.2 for the main preparations, Fig. 4.1 for the screening and Fig. 3.1 for 4D-Var.

a complex and multi-stranded flow of data through the IFS. Figure 1.1 illustrates how some of the data formats are used, and how the ODB-1 data is progressively augmented with additional information as the observations flow through the system. It also acts as a key to more detailed figures, found throughout this document, that summarise these steps in more detail (see figure caption).

However, the complexity of dataflow in the IFS makes it hard to summarise in a compact, well-structured way. Table 1.1 uncovers a little more of this complexity, and also summarises the structure used in this document. Chapter 2 covers observation preparations (defined in Sec. 1.1) which may take place before the main IFS, when it has historically been called preprocessing, or inside the main IFS code. The only place that has access to the atmospheric state is the main IFS, so we can identify the pre-processing stages as ‘state-independent’, as in Fig. 1.1. Note that the observation operators themselves are often responsible for processing and data selection decisions, so some preparations are described under Chapter 3, on the observation operators. Chapter 4 then covers the generic screening, i.e. the decisions on observation usage made inside the IFS (bearing in mind that some decisions have already been made during ‘preparation’ and are described in chapter 2 or 3). The generic screening is split into independent screening and dependent screening, such as redundancy checks and combined thinning. Finally, 4D-Var itself can update the status of observations, such as by downweighting outlying observations using VarQC.

1.4 OBSERVATION GROUPINGS IN THE IFS

Observations are associated with various different identifiers and grouping structures. Some illustrative though out-of-date tables are provided in Chapter 6; more current information is found in the ODB governance database at <http://apps.ecmwf.int/odbgov>. Externally-defined identifiers, such as WMO codes for satellite and instruments, and the BUFR type and subtype, are used in only a few parts of the IFS. Instead a set of internally-defined codes are used most widely. Of these the obstype, codetype and varno are documented by ODB governance. There are further, less well-known IFS groupings including the ‘SQNO’, the group tables and VarBC bias groups. Those groupings most relevant to the IFS will be discussed in the following sections. A long-term design goal must be to rationalise or eliminate some of

Table 1.1 *Observation processing stages in the IFS.*

Name	Format	Requirements
State-independent preparations (pre-processing) - Chapter 2		
Ingestion: in SAPP, before the IFS	BUFR, HDF	None
prepare_obs: e.g. simple thinning, reject obviously bad data, superobbing	BUFR	None
BUFR to ODB: not just conversion, but some processing too	BUFR, ODB-1	None
Preparations (including state dependent) - Chapter 2; Chapter 3		
setup: Many things including defining observation groupings like the satgrp table	ODB-1	None
defrun: Define some observation errors, QC paramers including VarQC and (for conventional data) Huber norm	ODB-1	None
Make CMA replacement: An ancient part of the code still needed for conventional data, but also important for satellite data, especially scatterometers, which retrieve wind from backscatter	ODB-1	FG
hretr: Mainly for satellite radiance data, including cloud and surface emissivity retrievals.	ODB-1	FG $H(x)$
observation operator: Many screening decisions and observation error assignments are handled when the observation operator is called, when LSCREEN=.TRUE.	ODB-1	FG $H(x)$
Independent screening - Chapter 4		
blocklist: utilisation decisions, e.g. excluding a particular sensor over sea-ice	ODB-1	FG $H(x)$
BgQC: background quality control, based on the first-guess departure and the observation errors	ODB-1	FG $H(x)$
Dependent screening - Chapter 4		
redundancy checks: (conventional observations)	ODB-1	FG $H(x)$ other obs
thinning: the 'best' remaining observation can be selected from among multiple candidates within a certain thinning box area	ODB-1	FG $H(x)$ other obs
4D-Var: observations may be removed by VarQC; analysis departures and other assimilation diagnostics are recorded	ODB-1	FG $H(x)$ other obs
Archiving		
Convert and archive: move to the ODB-2 format and archive to MARS	ODB-2	None

these groupings. Each grouping is associated with a considerable code and maintenance overhead, and moreover, badly-structured or overlapping groupings can be confusing.

1.4.1 MARS reporttype and group

The reporttype and group ID defines the structure of the ODB data archived in MARS. These are not intended to be used anywhere inside the IFS source code. Instead, the following internal IFS variables are used to drive different processing decisions. A particular reason not to use the REPORTTYPE in the IFS is that it is not extensible; every new satellite or sensor requires a new reporttype, so using it to drive decisions in the IFS could lead to difficult-to-maintain hardcoded lists of reporttypes in the source code.

1.4.2 IFS Obstype and codetype

There are currently 16 obstypes used in the IFS, and as such this is one of the coarsest possible groupings for the data. For example, obstype 7, 'SATEM' covers clear-sky radiances as well as satellite retrievals including some atmospheric composition observations. Although confusing, obstype is widely used to route observations through the more established parts of the IFS, in areas such as thinning, and to identify observations for special treatment. For example, in the horizontal interpolation from model to observation space (the GOMs) the choice of variable to interpolate, and how that interpolation is done, is dependent on the obstype. The diagnostic costfunction (the Jo table) is structured by obstype at its broadest level.

Codetype is sometimes used to provide a finer level of structure in areas that use obstypes. However, its meaning is not well defined. Codetypes are often used for conventional observations: they appear to be useful in distinguishing instrument types, for example radiosondes launched from land or ship ('Land TEMP' vs 'TEMP SHIP'), as well as coding practices ('Land TEMP' vs. 'BUFR Land TEMP'). For satellite data, they help distinguish radiances from retrievals, but this distinction would be better made at obstype level, or by using the varno (see later). Satellite codetypes are largely superfluous and could be deprecated in the future.

See <http://apps.ecmwf.int/odbgov> and Sec. 6.2 for more information.

1.4.3 SQNO - [DEPRECATED]

Not to be confused with the sequence number assigned to datums in the ODB, the 'SQNO' is a pair of indices that maps onto the obstype and codetype, used in a limited number of contexts in the IFS:

- Huber norm error definitions
- REPORT_EVENT2 and DATUM_EVENT2 categories
- Internal structure of the diagnostic Jo table

There is a one-to-one mapping between obstype and codetype and their SQNO equivalent, provided by the hardcoded routines and definitions in **CMOCTMAP**, **CMOCTMAP_INV** and **SUCOCTP**. The historical benefit of the SQNO is to provide a compressed index over the small set of codetypes being used for any obstype. Given the limited use cases, the annoyance of having to hardcode this information, and how easy it would be to eliminate it using modern code, this is deprecated and should be replaced in a future cycle.

1.4.4 IFS Physical variable: varno

The varno, or variable number, describes the type of variable being assimilated. This is the fundamental physical description, for example: '2m temperature in Kelvin'. As such the varno is well-defined and very useful. It is used throughout the screening and for selecting the appropriate observation operator. See <http://apps.ecmwf.int/odbgov> and Sec. 6.3 for more information.

1.4.5 Observation operator codes: NVAR

As further detailed in Sec. 6.4, conventional observation operators, and some other operators, have been given a number (NVAR) that maps directly onto a varno through `MAP_VARNO_TO_NVAR`. These codes are used to select conventional observations, to structure the observation error and QC settings for conventional observations in `DEFRUN` and as the third and lowest level of structure in the diagnostic Jo-table.

1.4.6 Satgrp_table and friends

For most satellite data, including limb radiances, ordinary radiances, AMVs, radio-occultation and scatterometers, the most useful groupings are by satellite ID and instrument. For example the instrument AMSU-A on satellite Metop-A forms one group. For satellite radiances, this structure is encoded in the `satgrp_table` by `YOMSATS` and `SURAD`. Similar tables are maintained for other satellite obtypes in `SUREO3`, `SULIMB`, etc. One very visible function in the IFS is to provide the relevant satellite groupings for the diagnostic Jo-table, and the descriptive names to go with that. A lot of other observation settings are encoded in the table, including those that help access the correct radiative transfer coefficient files for RTTOV. An important aspect of these tables is that they are dynamically created and only contain those satellite/instrument combinations available from the current observational data. Given the great similarities between the `satgrp_table` and those tables used by other satellite types, a design goal should be to provide a unified framework for creating these groups.

1.4.7 'Area' type - [DEPRECATED]

An overlapping concept to the `satgrp_table` is the misnamed 'area' type, maintained by `SUOBAREA`. It is only used for grouping satellite data when creating observation sets. Its original role related to the geographical area of conventional observations, long since superseded. Given its limited use, it should be eliminated in a future cycle.

1.4.8 VarBC bias groups

It's worth noting that VarBC maintains its own 'bias groups' distinct from any of the above grouping systems.

Chapter 2

Preparation of observations

Table of contents

- 2.1 Introduction**
- 2.2 State-independent preparations (preprocessing)**
 - 2.2.1 Clear-sky satellite radiances
 - 2.2.2 All-sky microwave radiances
 - 2.2.3 Ground-based radar precipitation composites
- 2.3 MKCMARPL tasks and functions**
 - 2.3.1 Basic observation processing setup
 - 2.3.2 Invoking, initializing and controlling the MKCMARPL
 - 2.3.3 MKCMARPL
 - 2.3.4 Basic observation handling routines
- 2.4 Variable conversions including retrievals**
 - 2.4.1 Conventional observations:
 - 2.4.2 Scatterometer winds
 - 2.4.3 Adjusted variables
 - 2.4.4 Other conversions
- 2.5 Observation errors**
 - 2.5.1 Conventional observations
 - 2.5.2 Satellite observations
 - 2.5.3 Ground-based radar precipitation composites
- 2.6 Background error estimates for screening**
- 2.7 Other preparations**
 - 2.7.1 Surface emissivity for radiances
 - 2.7.2 AMV height reassignment

2.1 INTRODUCTION

This chapter describes the processing required before a normalised first guess departure can be computed. This is best referred to as ‘observation preparation’ as it combines the ‘pre-processing’ tasks upstream of the main IFS code and the preparations that take place within the IFS during any ‘screening’ process in the trajectory (i.e. when `LSCREEN=.TRUE.`). Although the observation operator is part of this process, it is used more widely so it is left for the next chapter. Figure 1.1 and Table 1.1 show how this fits into the processing chain as a whole. Many of these preparations have historically been thought of as ‘screening’ simply because they have been performed during the screening run. This new documentation tries to make a clear distinction between ‘preparations’, which are described in this chapter, and ‘screening’, described in Chapter 4. Screening properly refers to what happens once the normalised first guess departures are available, as decisions on what goes into the assimilation are made at almost every stage of processing.

Historically, the upstream observation pre-processing (described here as ‘state-independent preparations’) have not been part of the IFS documentation. The aim is to bring this more fully into the documentation, but these tasks will not be covered in great detail. Fig. 2.1 illustrates these tasks but is far from comprehensive.

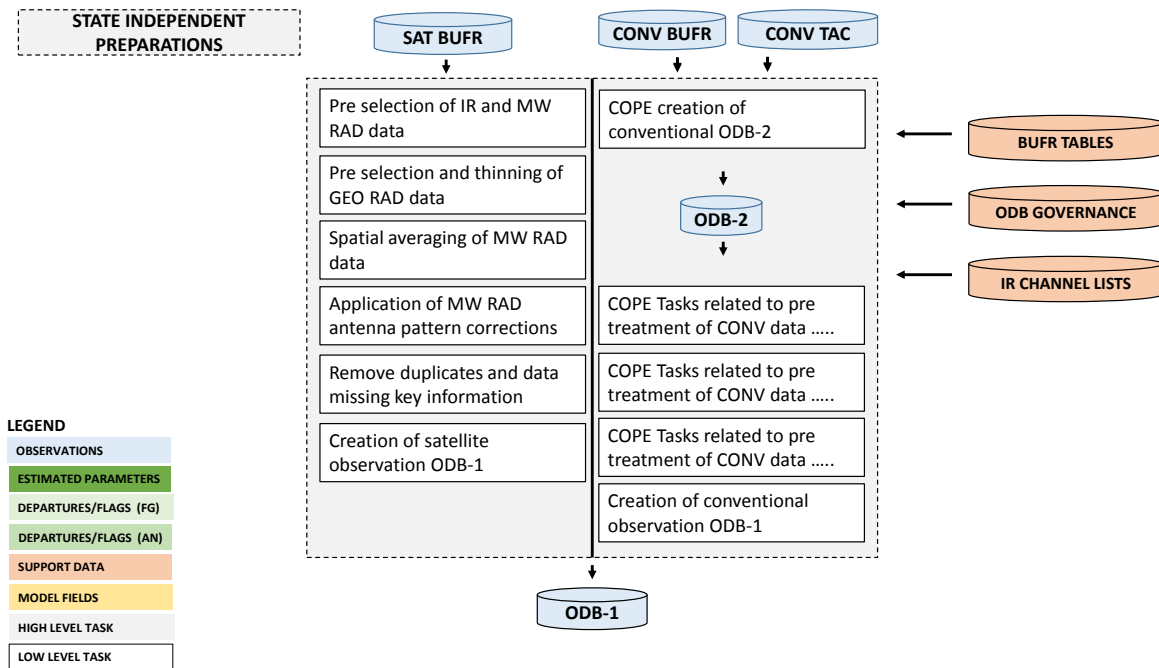


Figure 2.1 Simplified IFS flow diagram for state-independent observation preparations (alternatively known as preprocessing). The disc icons representing data stores and the rectangles representing processing stages. This is an illustrative, rather than comprehensive, description of the processing (NB. The COPE tasks mentioned are no longer used, the processing has reverted to MKCMARPL). See Fig. 1.1 for the wider observation processing context.

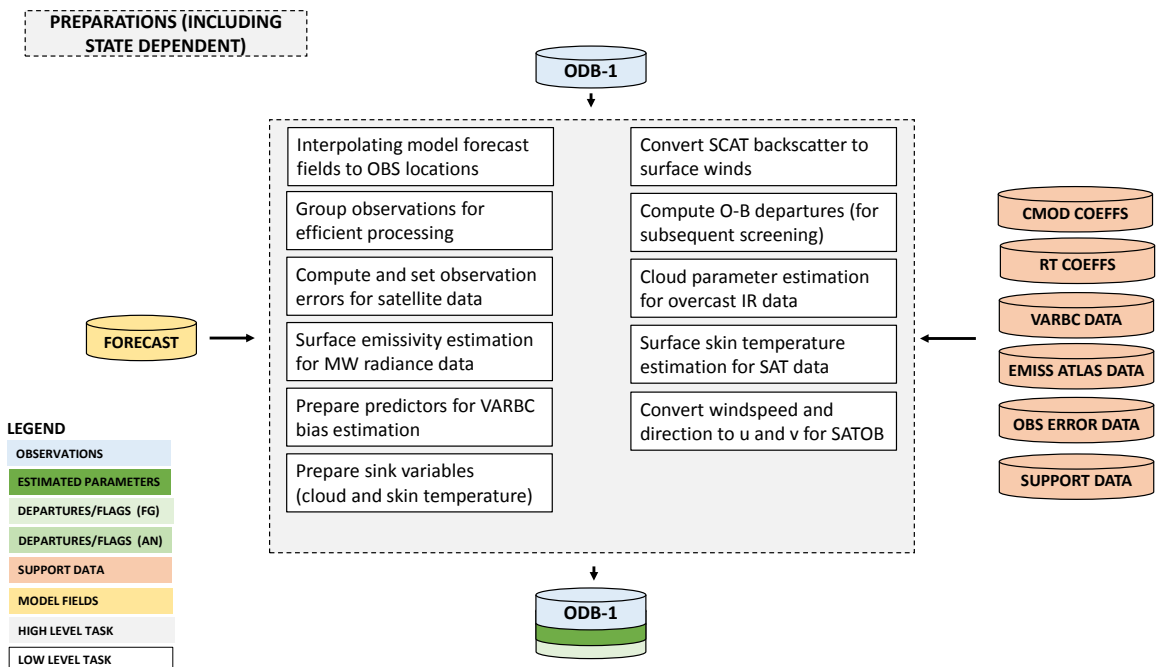


Figure 2.2 Simplified IFS flow diagram for observation preparations which may be state dependent: more broadly, processing which takes place inside the main IFS. The disc icons representing data stores and the rectangles representing processing stages. This is an illustrative, rather than comprehensive, description of the processing. See Fig. 1.1 for the wider observation processing context.

Figure 2.2 summarises broadly the preparations that take place inside the IFS and are mostly ‘state-dependent’. However, viewing the processing in this way is not always helpful. This chapter is mostly structured by the scientific procedure (e.g. parameter conversions, prior retrievals, observation error) rather than the technical location in which the science is achieved.

2.2 STATE-INDEPENDENT PREPARATIONS (PREPROCESSING)

All satellite data goes through the long-established path of BUFR preprocessing (the `PREPARE.OBS` tasks) and BUFR to ODB-2 conversion (the `BUFR2ODB` tasks). In this area, data undergo some rudimentary quality controls, e.g. a check for the observation format and position, and for the climatological limits. These jobs are multi-tasked running in parallel on multiple nodes. Several or all observation types can run synchronously.

2.2.1 Clear-sky satellite radiances

Radiance observations undergo a pre-screening process before being loaded into the ODB for input to the main IFS screening. Firstly, this is used to reduce the data volume and thus the computational burden of the main screening. Secondly, this rejects observations that fail to contain crucial header information and/or the correct number of channels that could potentially cause a computational run-time failure in the main screening. Observations in BUFR are decoded and checked inside the `SCREEN_1C_MOD` module called from the `BUFR.PRESCREEN` application where, additionally, data measured at particular scan lines and or scan positions may be removed to reduce the data volume (by setting `LINE_THIN`, `FOV_THIN` in the `PREPARE.PRESCREEN.CONFIG_*` configuration scripts). Observations which survive the checking and thinning process are then re-encoded in BUFR and supplied to the ODB loader. A key consideration for rejecting data in the pre-screening is that removed observations will NOT be passed through the IFS screening and thus will NOT accrue feedback quality information. Currently all pre-screening tasks are scalar (i.e. not parallel). However, for IASI (by far the largest data volume) the process is effectively parallelized by splitting the input BUFR file and launching multiple scalar tasks simultaneously.

2.2.2 All-sky microwave radiances

Observations are received in BUFR format and are pre-processed by `SCRIPTS/ECF_AN/bufr_prescreen.ecf`. The specific screening configuration for each instrument is set in `SCRIPTS/GEN/prepare_prescreen_config_*`. Observations part of the all-sky path are given codetype 215 (see `ODB/CMA2ODB/buf2cmat.new`).

For every all-sky instrument except AMSU-A, `SATRAD/PROGRAMS/bufr_grid_screen_parallel` is called to do superobbing. Observations are put onto a reduced Gaussian grid: currently N256 (40 km) for imagers and N200 (50 km) for sounders. The superob is a simple average of all observations whose geolocation falls within the Gaussian grid box, though a minimum number of observations are required per box to create a superob, a minimum that varies by instrument. In addition to greatly reducing data volume, this step brings the variable observation resolution of microwave instruments (as small as 10 km) to roughly the scale of clouds in the IFS model (Duncan *et al.*, 2024). The superobbing is done by computing the numerical mean of most BUFR fields, except the observation time, which comes from the last observation averaged, and non-continuous fields such as bit flags and azimuth angles.

Some microwave instruments feature non-co-located geolocations on the ground due to differently aligned feedhorns on the sensor platform. For these instruments (e.g. AWS, MWI and ICI) the superobbing step is performed on a per-feedhorn basis, followed by an additional step for matching the feedhorns together after superobbing to produce geolocated observation vectors. For such instruments, the superob geolocation is taken as the gridbox centre, whereas for instruments with aligned feedhorns (e.g. ATMS, MWS, MWHS2) the superob geolocation is the average of the longitude and latitude of constituent observations.

After this pre-processing, superobbed microwave radiances undergo thinning within the IFS prior to assimilation. Imagers are thinned to 110 km spacing, whereas sounders are thinned to 70 km spacing

for humidity channels and 100 km spacing for temperature channels. Because this thinning is within the IFS, the full set of superobbed radiances are available after the assimilation for diagnostic analysis. AMSU-A is a special case, with thinning performed outside the IFS at N80 resolution (125 km) and in 30 minute time slots considering all AMSU-A sensors together.

2.2.3 Ground-based radar precipitation composites

Prior to their assimilation, the original hourly 4-km NCEP Stage IV surface precipitation composites (based on NEXRAD ground-based radars and (some) rain gauges over the United States; Lopez, 2011) are averaged over the model grid at trajectory resolution to reduce possible representativeness issues. In addition, only observations that are located east of 105°W are selected. This is to avoid the likely degradation of ground-radar-based measurements over the Rocky Mountains due to beam blocking, precipitation enhancement or the need for a proper correction of vertical profiles of reflectivity at higher radar beam angles. For similar reasons, observations located over high ($z_{\text{orog}} > 1500$ m) or rugged terrain ($\sigma_{\text{orog}} > 100$ m) are rejected. The observations are then accumulated over the period NPRACCL, currently set to 21600 seconds in namelist NAEPHY. All this pre-screening is performed by routine `BUFR_SCREEN_NEXRAD`. Pre-processed observations are then converted to $\log(\text{RR}[\text{mm/h}]+1)$ space in routine `B2O_CONVERT_RAIN_RATES`.

2.3 MKCMARPL TASKS AND FUNCTIONS

2.3.1 Basic observation processing setup

In order to perform the observation processing functions, a number of basic observation processing setups are carried out at the very beginning of initialising the IFS. This is done by calling several routines in addition to all other routines needed to setup the IFS.

- Program `MASTER` calls `CNT0` which in turn calls `SU0YOMA`.
- `SU0YOMA` calls (among other routines) `SUOAF` from which `SUCMOCTP`, `SUEVENTS`, `SUCODES`, `SUFLTXT` and `SUCMA` are called. `SUCMOCTP` defines the ODB observation types and code types, and `SUEVENTS`, `SUCODES` and `SUFLTXT` define analysis events, various codes used and flags naming conventions.
- `SUCMA` calls `SUCMAF` which then calls several subroutines: `SUCMAD1`, `SUCMAD2`, `SUCMAHFP`, `SUCMAHOP`, `SUCMBDFP` and `SUCMBDTP`. These routines define the structure of ODB Data Descriptor Records (DDR) as well as the ODB packing patterns (bit structure) employed for header and body respectively.

2.3.2 Invoking, initializing and controlling the MKCMARPL

The MKCMARPL run is initiated by the `MKCMARPL` subroutine. This routine is only invoked in the SCREENING run of the IFS. It is called, together with some of its additional setup routines via subroutine `SUOBS`. The additional setup routines called at this level are: `SUANCT`, `DEFRUN`, `SULIM`, `SULEVLAY`, `SUSATRET`, `SUVNMB`, `SUSCRE0`, `SUOBSORT`, `SETCOM`, `DEPERERR`, `SUERRORS`, `INIERSCA`.

- `MKCMARPL` is namelist driven and in `DEFRUN` a logical variable `LMKCMARPL` is defined. By default `LMKCMARPL = .TRUE.` but it can be overwritten via namelist `NAMOBS`. Furthermore, many other parameters and switches are defined in `DEFRUN` and some of them can also be overwritten via namelists.
- `SUANCT` and `SULIM` define some additional analysis constants and limits.
- `SULEVLAY` and `SUSATRET` define analysis related level/layer and satellite retrieval parameters, respectively.
- `SUVNMB` declares variable numbers.
- `SUSCRE0`, `SUOBSORT` and `SETCOM` define flag limits, identify ambiguous moving platforms, initialise observation sorting, and provide some general observation common variables.
- `DEPERERR` and `SU_ERRORS` deal with observation error statistics definitions. `SU_ERRORS` calls `SUOBSERR` to define prescribed observation errors.

- **INIERSCA** deals with initialising SCATT processing.

The next step is to find out if it is a SCREENING run and if so to check if it is a MKCMARPL run as well. In the case of a MKCMARPL run all aspects of the observation processing before the screening are dealt with by calling **MKCMARPL** (more about it in [Subsection 2.3.3](#)). After **MKCMARPL** has finished there are several ways to proceed. These depend on the status of LMKCMARPLO and LRPLSWAPOUT logical switches (NAMOBS namelist). If LRPLSWAPOUT = .TRUE. the ODB is swapped out and if LMKCMARPLO = .TRUE. the ODB is written out and the run terminated. Both of these options are not normally used and their use is for diagnostics/debugging purposes. Once the **MKCMARPL** work has been completed the remainder of SUOBS will execute as before. Thus, calls to **WRITEOBA**, **WINDAUX**, **OBATABS**, **SUAMV**, **SURAD**, **SULIMB**, **SUOBAREA**, **MKGLOBSTAB** and **SUREO3** are issued.

In the context of operational running, the **MKCMARPL** related switches are set:

```
LMKCMARPL = .TRUE.  LRPLSWAPOUT = .FALSE.  LMKCAMRPLO = .FALSE.
```

2.3.3 MKCMARPL

The main purpose of **MKCMARPL** is to control the IFS observation pre-processing. Observation pre-processing at this stage is done in groups of observations. At the moment there are six groups: CONVENTIONAL, SATOB, TOVS/RTOVS, SCATT, LEVEL1C/GEOSS and OZONE observations. For each group a separate subroutine is called: **CONVENTIONAL_OB**, **SATOB_OB**, **TOVSRTOVBS_OB**, **SCAT_OB**, **LEVEL1CGOES_OB** and **OZONE_OB**. These routines are just cover or hat routines for the actual work to be carried out underneath. However, **TOVSRTOVBS_OB** is currently not called because it is obsolete.

Each cover routine would call the ODB to get the observations it wants to process. This is done by calling the ODB **GETDB** subroutine. As the observations are brought, in one or more worker routines would be called to perform the observation processing functions. Once the worker routines have finished the control is handed back to the cover routine. The next step in the cover routine is to return observations back to the ODB database. This is done by calling the ODB **PUTDB** routine. In some of these cover routines several calls to **GETDB/PUTDB** might be issued. This is because there may be sufficient differences between similar data to justify a slightly different approach in their pre-processing. For example under the **CONVENTIONAL_OB** routine there are two calls to a **GETDB** and **PUTDB** pair. The first call deals with all conventional observations except SATEMS; the second call deals with the SATEMS. As indicated earlier, between each **GETDB** and **PUTDB** a number of observations type or code type designed worker routines are called.

- **CONVENTIONAL_OB** calls the following worker routines: **SYNOPIN**, **AIREPIN**, **DRIBUIN**, **TEMPIN**, **PILOTIN**, **EWPRFIN**, **AWPRFIN** and **METARSIN**. A worker routine name indicates which observations it is dealing with.
- **SATOB_OB** calls **SATOBIN** and **SATAMIN**.
- **SCAT_OB** calls **ERS1IN**, **NSCATIN**, **ASCATIN** and **QSCATIN**.
- **LEVEL1CGEOS_OB** calls **RAD1CIN** and **GOESRIN**.
- **OZONE_OB** calls only **REO3SIN**.

2.3.4 Basic observation handling routines

The observation pre-processing worker routines referred to in [Subsection 2.3.3](#), names of which always end with "IN", are the basic observation handling routines. They all follow more or less the same logic. As an example consider **AIREPIN** which deals with AIREP observations.

The first thing which is done is to define the instrument specification (**OBINSTP**) followed by preliminary quality control check both at the report level (**PRLMCHK**) as well as at the data level (**GETSETE** and **AIREPBE**).

- **PRLMCHK** calls **REPSEL** and **TIMDIF** to do report selection according to preset criteria and to find out time difference between analysis time and the actual observation time, respectively.

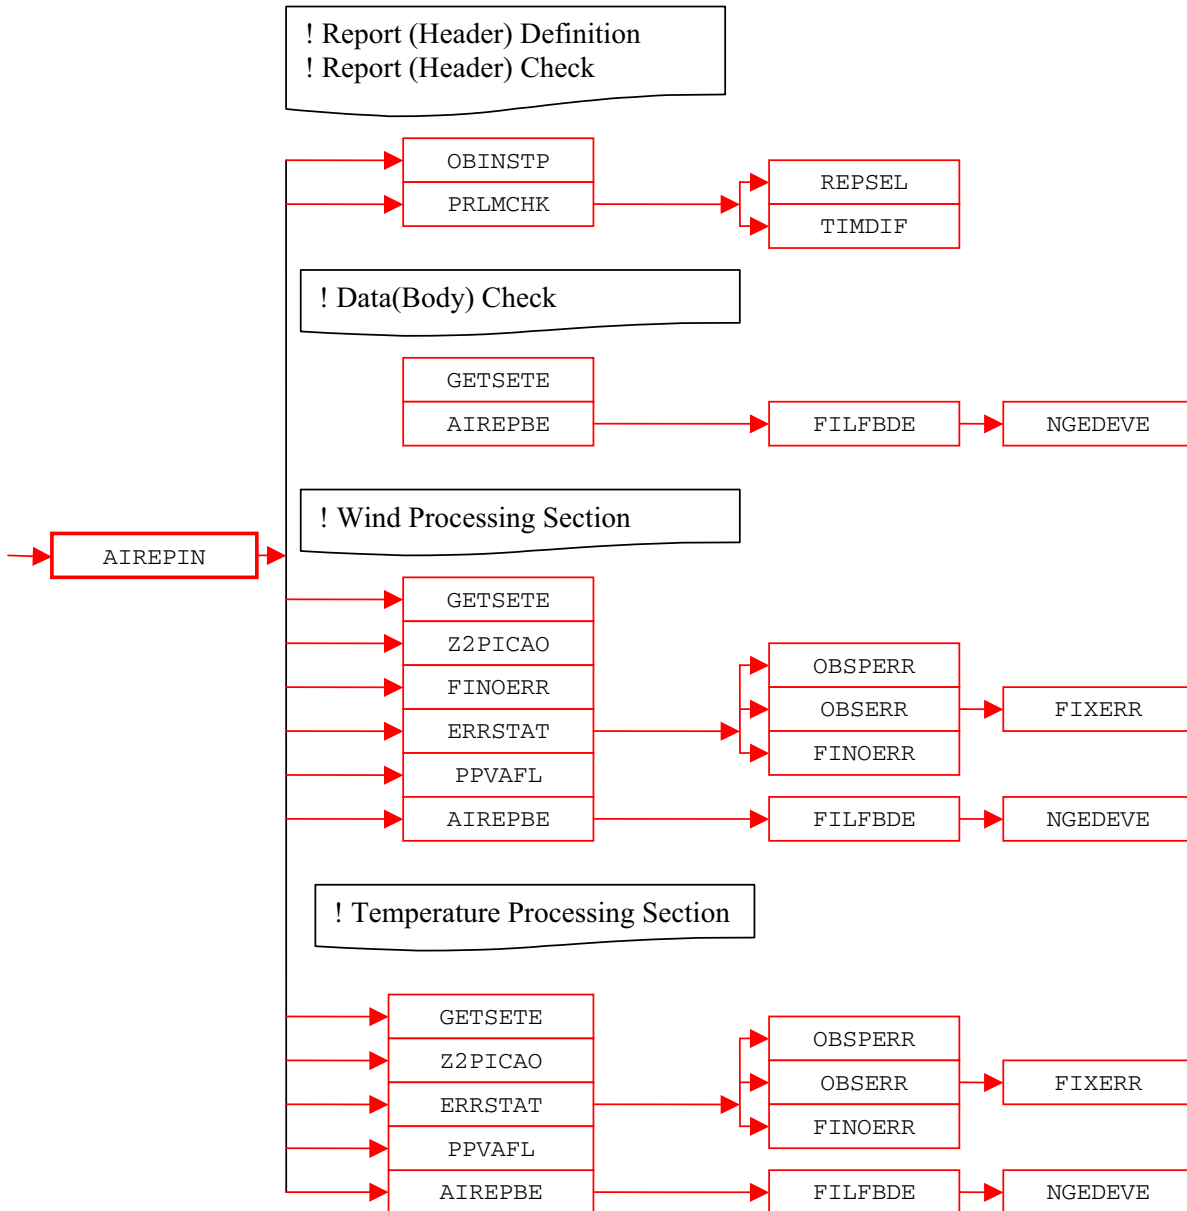


Figure 2.3 Simplified IFS observation pre-processing flow diagram, with AIREPIN as an example. Colour coding scheme: routines in red boxes perform observation pre-processing.

- **GETSETE** makes a local copy of a given observation variable and its related parameters from an ODB supplied array.
- After updating the local copy, **AIREPBE** is called to return the updated local copy back to the ODB supplied array.

The preliminary quality control at the report level consists of making sure that observation position, date and time are reasonable. Furthermore, as there is a possibility of excluding certain observations via the NAMOBS namelist, a check is made of whether the observation is actually wanted at this stage. Once the report level check is passed attention is turned to the data itself. Each datum is checked against predefined list of expected data. If not in the list, datum is rejected and a warning message issued. At this stage it is also ensured that missing indicators used are unique.

After the preliminary phase attention is turned to getting data in the right form and shape for further usage. Thus, in the case of an AIREP observation, this is done in sections of available variables: wind and temperature.

- (i) **Wind.** There are four wind variables: wind direction (*DDD*), wind force (*FFF*), *u* and *v* components. For each of these variables the first thing which is done is to get a local copy of it together with its related parameters from an ODB supplied array (**GETSETE**). Once a variable is made available locally a check is made to ensure that the vertical coordinate is pressure; if instead of pressure a flight level is supplied it is converted into pressure by assuming a standard ICAO atmosphere (**Z2PICAO**). If the variable in question is either *u* or *v*, then *DDD* and *FFF* are converted into *u* and *v* wind components. Furthermore, for each of the four variables appropriate observation error statistics are assigned (**ERRSTAT**, **FINOERR**). Also, if any flags are set at this stage an appropriate word in the local copy is updated (**PPVAFL**). Finally, an updated local copy of an observed quantity and its related parameters are returned back into the ODB (**AIREPBE**).
- (ii) **Temperature.** In the case of temperature only one observed variable is dealt with. The pattern of making a local copy (**GETSETE**), ensuring that pressure is the vertical coordinate (**Z2PICAO**), assigning the observation error statistics (**ERRSTAT**), updating flags (**PPVAFL**) and returning an updated local copy back to the ODB (**AIREPBE**) is repeated.

As just mentioned **ERRSTAT** deals with assigning observation errors for a given observation variable. **ERRSTAT** calls **OBSERR** which in turn calls **FIXERR** to assign prescribed observation error. It is worth mentioning that observation errors themselves are already predefined at an earlier stage (**SU_ERRORS**).

The pattern of activities outlined for **AIREPIN** is repeated more or less in the other worker routines. However, the **SYNOPIN** routine is first split further into **SHIPIN**, **METARIN**, **PGPSIN** and **LANSYIN**. This is because SHIP, METAR, GPS and SYNOP LAND observations are sufficiently different to justify a separate worker routine. Furthermore, **LANSYIN** is somewhat more complicated than **AIREPIN**. One of the reasons for this is that we have to distinguish between low and high level stations.

2.4 VARIABLE CONVERSIONS INCLUDING RETRIEVALS

Most observations are used in their original form as this is usually the most optimal use of the data. However, some are transformed, either with a change of variable, but also there may be a retrieval from satellite data if they are independent from the background model fields. The original variables may be kept with the derived ones so that first guess departures can be assigned for both.

2.4.1 Conventional observations:

Variables which are transformed for further use by the analysis are as follows.

- (i) Wind direction (*DDD*) and force (*FFF*) are transformed into wind components (*u* and *v*) for SYNOP, AIREP, SATOB, BUOY (aka DRIBU), TEMP and PILOT observations.
- (ii) Temperature (*T*) and dew point (*Td*) are transformed into relative humidity (*RH*) for SYNOP and TEMP observations, with a further transformation of the *RH* into specific humidity (*Q*) for TEMP observations.

The wind components are worked out as

$$u = -FFF \sin\left(DDD \frac{\pi}{180}\right)$$

$$v = -FFF \cos\left(DDD \frac{\pi}{180}\right)$$

The *RH* is derived using

$$RH = \frac{F(Td)}{F(T)}$$

where function F of either T or Td is expressed as

$$F(T) = a \frac{R_{\text{dry}}}{R_{\text{vap}}} e^{b \frac{T-T_0}{T-c}}$$

where $T_0 = 273.16 \text{ K}$, $a = 611.21$, $b = 17.502$, $c = 32.19$, $R_{\text{dry}} = 287.0597$ and $R_{\text{vap}} = 461.5250$ are constants.

The specific humidity Q is worked out by using

$$Q = RH \frac{A}{1 - RH \left(\frac{R_{\text{vap}}}{R_{\text{dry}}} - 1 \right) A}$$

with function A is expressed as

$$A = \min \left[0.5, \frac{F(T)}{P} \right]$$

where P is pressure. Q is assigned in the **RH2Q** subroutine.

These transformations are done in MKCMARPL (Sec. 2.3).

For SYNOP, SHIP and BUOY reports the usual humidity variable reported is dewpoint temperature and for many years this was the only one extracted. However some (particularly in cold regions) report relative humidity and from 49R1 this is also extracted. If both humidity variables are reported then dewpoint temperature is used.

2.4.2 Scatterometer winds

For ASCAT and possibly some other scatterometers, the retrieval from backscatter to wind is done in 'Make CMA replacement' process (Sec. 2.3).

Dedicated observation operators exist for scatterometer ambiguous surface winds (Stoffelen and Anderson, 1997; Isaksen and Janssen, 2004). Backscatters (σ^0 's) are transformed into several pairs of ambiguous wind components (u and v); this actually involves a retrieval according to some model function describing the relationship between winds and σ^0 's and requires a fair bit of computational work. The screening of scatterometer data involves the conversion of the backscatter measurements acquired by the instrument (triplets for ERS and ASCAT and quadruplets for NSCAT and QuikSCAT) into ambiguous u and v wind components that will actually be assimilated into the IFS (see Section 3.2.9). The (empirical) relation between wind and backscatter is described by a geophysical model function (GMF). Although in principle inverted wind components are provided as a level 2 product, at ECMWF the wind inversion is performed in house. In this way any drifts in backscatter levels can be corrected in a direct manner.

Data from the AMI instrument on ERS-2 was used from June 1996 (with an interruption from January 2001 till March 2004) until September 2011, data from the SeaWinds instrument on-board QuikSCAT was used from January 2002 until November 2009 (when QuikSCAT failed), data from ASCAT on MetOp have been assimilated from June 2007 onwards. Oceansat-2 data was used from December 2012 to February 2014 and data from HSCAT on HY-2B has been assimilated since June 2021. Data from NSCAT and RapidSCAT have never been used in an operational setup, although offline assimilation experiments have been performed. From November 2010 onwards scatterometer data is assimilated as equivalent-neutral 10-metre wind, rather than (real) 10-metre wind, since the former model winds are closer related to scatterometer observations.

(a) Wind retrieval

Since geometry and measurement principle of ERS and ASCAT are alike, data from these instruments is processed in a similar way. The procedure for wind inversion closely follows the wind retrieval and ambiguity removal scheme originally developed for the ERS-1 scatterometer (Stoffelen and Anderson, 1997), though the original geophysical model function CMOD4 has been replaced by CMOD5 (Hersbach

et al., 2007) in March 2004, by CMOD5.4 in June 2007, and by CMOD5.n (Hersbach, 2010b) in November 2010, after which scatterometer winds are assimilated as equivalent-neutral wind (Hersbach, 2010a).

For QuikSCAT the task of wind inversion is performed in the pre-screening (PRESCAT). Data are like ERS and ASCAT, provided at a resolution of 25 km. Rather than data thinning (see Subsection 4.5.6), for QuikSCAT a 50 km product is created which contains information about backscatter from the four underlying original sub-cells. The weight of the scatterometer cost function (defined in routine HJO) of each 50 km wind vector cell is reduced by a factor four, which effectively mimics the assimilation of a 100 km product. It is the re-sampled 50 km product that is stored in ODB. Original backscatter observations at 25 km are not available within the assimilation.

In general, the wind retrieval is performed by minimizing the distance between observed backscatter values σ_{oi}^0 and modelled backscatter values σ_{mi}^0 , given by

$$D(u) = \sum_i^n \frac{[(\sigma_{oi}^0)^p - \sigma_{mi}^0(u)^p]^2}{k_p \left[\sum_j^n \sigma_{mj}^0(u)^p \right]^2} \quad (2.1)$$

For ERS and ASCAT data, the sum is over triplets, while for QuikSCAT the sum may extend to 16 values (four 25 km sub-cells with each four observations). The quantity p is equal to unity for NSCAT and QuikSCAT. For ERS and ASCAT data, a value of $p = 0.625$ was introduced because it makes the underlying GMF more harmonic, which helps to avoid direction-trapping effects (Stoffelen and Anderson, 1997). The noise to signal ratio k_p provides an estimate for the relative accuracy of the observations.

The simulation of σ_m^0 is for ERS and ASCAT data based on the CMOD5.4 model function. For NSCAT data the NSCAT-2 GMF has been utilized. For QuikSCAT data, the choice of GMF is handled by a logical switch LQTABLE. By default LQTABLE = .TRUE. and the QSCAT-1 model function is used, otherwise, modelled backscatter values are based on the NSCAT-2 GMF. The minimization is achieved using a tabular form of the GMF, giving the value of the backscatter coefficient for wind speeds, direction and incidence angles discretized with, for ERS and ASCAT data, steps of 0.5 ms^{-1} , 5° and 1° , respectively. For NSCAT and QuikSCAT data the corresponding values are 0.2 ms^{-1} , 2.5° and 1° . ERS and ASCAT use the same table, which is read in the initialisation subroutine INIERSCA. For QuikSCAT, inversion takes place in the QSCAT25TO50KM program in the PRESCAT task.

(b) Quality control

The wind inversion involves some quality control. For ERS (ERS1IF), k_p must for each antenna be below 10%, and a missing packet number must be less than 10 to ensure that enough individual backscatter measurements have been averaged for estimating the value.

For ASCAT (ASCATIF) a in the product provided land fraction must be zero for each backscatter measurement. No restriction on k_p is imposed, other than that values should be non missing. It is checked whether two other provided quality flags ('sigma0 usability' and 'kp quality') have acceptable values. However, no quality control decisions are made on these two indicators for the moment, since so far, they have not been fully calibrated and validated by EUMETSAT.

For QuikSCAT, from 38 across-track 50 km cells, the outer 4 at either side of the swath are, due to their known reduced quality rejected. In addition, for QuikSCAT, it is verified whether inverted winds are well-defined, i.e. whether minima $D(u)$ are sufficiently sharp. In practise this is mainly an issue for cells in the central part of the swath. Data is rejected when the angle between the most likely solution and its most anti-parallel one is less than 135° (routine SCAQC).

After wind inversion, a further check is done on the backscatter residual associated to the rank-1 solution (also called 'distance to the cone'). This misfit contains both the effects of instrumental noise and of GMF errors. Locally, these errors can become large when the measurements are affected by geophysical parameters not taken into account by the GMF, such as sea-state or intense rainfall. For ERS, a triplet is rejected when the cone distance exceeds a threshold of three times its expected value. For QuikSCAT and ASCAT data such a test is not performed.

In addition to a distance-to-cone test on single observations, a similar test is performed for averages for data within certain time slots. If these averages exceed certain values, all data within the considered time slot is suspected to be affected by an instrument anomaly, since geophysical fluctuations are expected to be averaged out when grouping together large numbers of data points. For ERS and ASCAT, cell-wise averages are calculated for the default screening time slot (30 minutes) in the IFS routine `SCAQC`, and its rejection threshold (1.5 times average values) are defined in the IFS routine `SUFGLIM`. For QuikSCAT averages are considered over six-hourly data files and are evaluated in the pre-screening (`DCONE.OC`), using a threshold of 1.45 for any of cells between 5 and 34.

(c) *Rain contamination*

Thanks to the usage of C-band frequency, rain contamination is mild for ERS and ASCAT. For QuikSCAT and NSCAT, which operate in Ku band, rain contamination is a serious issue.

For QuikSCAT the check on rain contamination occurs in the pre-screening and is imposed on the original 25 km observations. Any 25 km rejected cell is not used in the determination of the 50 km wind product. When more than one 25 km sub-cell is rejected, the entire 50 km product is rejected (decision made in `SCAQC`).

Since February 2000, the BUFR product provides a rain flag. This flag, which was developed by NASA/JPL, is based on a multidimensional histogram (MUDH) incorporating various quantities that may be used for the detection of rain (Huddleston and Stiles, 2000). Examples of such parameters are `mp_rain_probability` (an empirically determined estimate for the probability of a columnar rain rate larger than $2 \text{ m}^2 \text{ hr}^{-1}$; typically values larger than 0.1 indicate rain contamination) and `nof_rain_index` (a rescaled normalized objective function – values larger than 20 give a proxy for rain). Since at the time of implementation, the quality of the JPL rain flag had not been fully confirmed, an alternative (more aggressive) flag was established in house. Based on a study in which QuikSCAT winds were compared to collocated ECMWF first guess winds, a quality flag was introduced. It is given by

$$L_{\text{rain}} = \{ \text{nof_rain_index} + 200 \text{ mp_rain_probability} > 30 \}.$$

Both `mp_rain_probability` and `nof_rain_index` are provided in the original 25 km BUFR product (for details see Leidner *et al.*, 2000). When one of these quantities is missing, the above mentioned condition for the remaining quantity is used.

(d) *Bias corrections*

For ASCAT and ERS, bias corrections are applied, both in terms of backscatter (before wind inversion) and wind speed (after inversion), particularly to compensate for any change in the instrumental calibration and to ensure consistency between the retrieved and model winds. The backscatter and wind-speed bias corrections are defined by dedicated files read in the initialization subroutine `INIERSCA`. Files are in principle model-cycle and date dependent. Currently for ERS-2, the appropriate files have no effect (i.e. containing only unity correction factors and zeros), since the CMOD5.4 GMF was tuned on ERS-2. For ASCAT, though, the usage of bias corrections is essential, since the backscatter product for this instrument has been calibrated differently from ERS. The bias correction file for backscatter has been updated every time a change in the calibration of ASCAT was imposed by EUMETSAT.

For QuikSCAT data no bias corrections in σ^0 space is applied, though, wind-bias corrections are made. This also takes place in the pre-screening. Corrections are performed in three steps. First of all, wind speeds are slightly reduced according to:

$$v' = 0.2 + 0.96 v.$$

Where v is the wind speed as obtained from inversion (2.1) The addition of 0.2 ms^{-1} is used in the operational configuration, where scatterometer data is assimilated as equivalent neutral wind. In case this is not desired (expressed by `LSCATT_NEUTRAL=false`) only the rescaling factor of 0.96 is used. It was observed that the residual bias between QuikSCAT winds and ECMWF first guess winds depends on the value of `mp_rain_probability`. The motivation is that, for higher amounts of precipitation, a larger

part of the total backscatter is induced by rain, leaving a smaller part for the wind signal. The correction applied is

$$v'' = v' - 20 \langle mp_rain_probability \rangle,$$

where $\langle \rangle$ denotes the average value over the 25 km sub-cells that were taken into account in the inversion (i.e. over rain-free sub-cells). The maximum allowed correction is 2.5 ms^{-1} , which is seldom reached. Finally, for strong winds, QuikSCAT winds were found to be quite higher than their ECMWF first guess counterparts. In order to accommodate this, for winds stronger than 19 ms^{-1} the following correction is applied:

$$v''' = v'' - 0.2(v'' - 19.0).$$

2.4.3 Adjusted variables

For radiosonde and synoptic stations the official WMO source of metadata is OSCAR/Surface (<https://oscar.wmo.int/surface/>). Station position metadata are also available in BUFR. Investigations at ECMWF suggest that where they differ significantly OSCAR/Surface is more likely to be correct (but not always). In the SAPP system the latitude, longitude and height from OSCAR/Surface are appended to the BUFR reports. If the values are in good agreement then the BUFR positions are used, if they differ by more than 0.1 degrees latitude/longitude or 5 m in height then the OSCAR values are used. (B2O_STATION_POSITION_TOLERANCE and B2O_STATION_HEIGHT_TOLERANCE can be modified in prepIFS.) For station height we use barometer height (if available) because pressure is more sensitive to the height than other variables.

Radiosonde temperature and humidity profiles are bias corrected based on [Agusti-Panareda et al. \(2009\)](#). The corrections are based on accumulated monitoring statistics (updated monthly) and are a function of radiosonde type (SONDE.TYPE) and solar elevation angle. A reference subset of radiosonde data is chosen: originally nighttime Vaisala RS92 but this was changed to be nighttime Vaisala RS92+RS41 during the transition to the newer RS41 radiosonde. The humidity bias correction is switched off for RS41. There is an implicit assumption that background biases are similar at the different radiosonde sites. In general more modern radiosonde types have smaller biases. This is just as well because the bias correction program does not work properly for radiosonde types over 99, that became possible with the move to BUFR reporting, so these are not bias corrected. Tests switching off radiosonde bias correction show that it is still marginally useful. The bias correction is used to change the 'observed values' (in contrast to VarBC where OBSVALUE is unchanged and the correction is applied within the observation operator). The humidity correction is applied to relative humidity before the conversion to specific humidity.

A bias correction is also applied to surface pressure. Up to 49r1 operations used the method of [Vasiljevic \(2006\)](#) although ERA5 used VarBC. From 49r2/50r1 a simpler but more robust Python-based correction (`scripts/gen/psbiasup.py`) is used both in ERA6 and operations. This is based on monitoring statistics over the last 14 days updated every 12 hours. Essentially a robust mean O-B is calculated for each station. This estimate of the bias is used to correct new observations if a) the sample was large enough (over 20 points contributed), b) the robust standard deviation was less than 1.5 hPa and c) the magnitude of the correction is greater than a threshold. For buoys the threshold is 0.9 hPa, for other surface stations the threshold is 1.2 hPa for latitudes less than 40 degrees increasing linearly to 1.7 hPa at the poles. The larger threshold at high latitudes reflects the fact that local model biases can be larger there, especially in winter. Under the previous scheme most research experiments used the operational bias corrections rather than calculating their own. Under the new scheme most research experiments will calculate their own (starting from a 14 day archive of the experiment used for the initial conditions).

Data from Boeing 787 aircraft in AIREP or ADS-C reports are subject to a wind direction error when the wind is southerly. A correction, changing the sign of v , is applied (in SWITCH_VWIND if LAIRSWITCHV is true) if it gives a significantly better fit to the background. The test is $|v_o + v_b| < \min(5, |v_o - v_b| - 8)$ with values in m/s. There is no check on aircraft type because, unfortunately, the metadata are patchy. A few more details are given in [Pauley and Ingleby \(2022\)](#).

2.4.4 Other conversions

It is worth mentioning the superobbing of all-sky microwave radiances from their raw resolution into a superob more representative of the cloud and precipitation scales in the forecast model. This is done in the BUFR stages of pre-processing.

2.5 OBSERVATION ERRORS

The diagonal part of the observation errors are written, at some point during the observation processing, into the ODB column MDBFOE for use by the data assimilation algorithm. However, the source of these error estimates is varied.

2.5.1 Conventional observations

In the past three types of observation errors were dealt with at the observation pre-processing level: persistence observation error, prescribed observation error, and the combination of the two above called the final observation error. Persistence observation error is no longer used. The prescribed observation error is described below, as well as the additional error inflation for dropsondes, and T2m data which is done within the data assimilation system, as a function of the FG departure.

(a) Prescribed observation errors

Prescribed observational errors have been derived by statistical evaluation of the performance of the observing systems, as components of the assimilation system, over a long period of operational use. Currently, observational errors are defined for each observation type that carries the following quantities.

- (i) Wind components.
- (ii) Height.
- (iii) Temperature.
- (iv) Humidity.

As can be seen from the tables of prescribed observation errors, they are defined at standard pressure levels but the ones used are interpolated to the observed pressures. The interpolation is such that the observation error is kept constant below the lowest and above the highest levels, whereas in between it is interpolated linearly in $\ln p$. Several subroutines are used for working out the prescribed observation error: **SUOBSERR**, **OBSERR**, **FIXERR**, **THIOERR** and **PWCOERR**.

- **SUOBSERR** defines observation errors for standard pressure levels. For radiosondes six separate error profiles are used for different radiosonde types (SONDE_TYPE) based on statistics like those in [Ingleby \(2017\)](#).
- **OBSERR** and **FIXERR** calculate the actual values.
- **THIOERR** and **PWCOERR** are two specialised subroutines to deal with thickness and PWC errors.

(b) Derived observation errors

Specific humidity observation error, Q_{err} , is a function of RH , RH_{err} , P , P_{err} , T and T_{err} , and formally can be expressed as

$$Q_{\text{err}} = Q_{\text{err}}(RH, RH_{\text{err}}, P, P_{\text{err}}, T, T_{\text{err}})$$

or

$$Q_{\text{err}} = RH_{\text{err}} F_1(RH, P, T)$$

where function F_1 is given by

$$F_1(RH, T, P) = \frac{A}{\left[1 - RH \left(\frac{R_{\text{vap}}}{R_{\text{dry}}} - 1\right) A\right]^2}$$

Subroutine **RH2Q** is used to evaluate Q_{err} .

Surface pressure observation error $P_{s_{\text{err}}}$ is derived by multiplying the height observation error Z_{err} by a constant:

$$P_{s_{\text{err}}} = 1.225 Z_{\text{err}}$$

However, the $P_{s_{\text{err}}}$ may be reduced if the pressure tendency correction is applied. For non-SHIP data the reduction factor is 4, whereas for SHIP data the reduction factor is either 2 or 4, depending on if the P_t is adjusted for SHIP movement or not.

The thickness observation error (DZ_{err}) is derived from Z_{err} .

(c) *Final (combined) observation error*

In addition to the prescribed observation and persistence errors, the so called final observation error is assigned at this stage too. This is simply a combination of observation and persistence errors given by

$$F_{\text{OE}} = \sqrt{O_{\text{E}}^2 + P_{\text{E}}^2}$$

where F_{OE} , O_{E} and P_{E} are final, prescribed and persistence observation errors, respectively. The subroutine used for this purpose is **FINOERR**.

(d) *Additional error inflation for dropsondes*

Dropsonde errors can be inflated further as a function of the first guess departure d , in order to prevent problems in tropical cyclones, where the representation error can be large. This is done in **FGWND**. A brief summary is given here; more details can be found in [Bonavita et al. \(2017\)](#). If the 'final' observation error is F_{OE} then a new and more final observation error F_{OE2} is computed by adding a component for representation error R_{O} :

$$F_{\text{OE2}}^2 = F_{\text{OE}}^2 + R_{\text{O}}^2$$

where for $d^2 \leq F_{\text{OE}}^2 + B_{\text{O}}^2$

$$R_{\text{O}}^2 = 0$$

and for $d^2 > F_{\text{OE}}^2 + B_{\text{O}}^2$

$$R_{\text{O}}^2 = d^2 - F_{\text{OE}}^2 - B_{\text{O}}^2$$

Here, B_{O} represents an estimate of background error in observation space.

(e) *Error inflation for T2m*

In a similar way T2m errors can be inflated (in **FGCHK** if `LT2.TAILS` is true) as a function of the first guess departure d . This was found necessary to avoid excessive analysis increments particularly in high-latitude, winter conditions. No inflation is performed if $d^2 \leq F_{\text{OE}}^2 + B_{\text{O}}^2$. Otherwise

$$F_{\text{OE2}} = F_{\text{OE}} \sqrt{d^2 / (F_{\text{OE}}^2 + B_{\text{O}}^2)}$$

2.5.2 Satellite observations

(a) *Clear-sky radiances*

The setup routine **DEFRUN** sets up initial defaults for clear-sky radiance observation errors, set by sensor and channel in the variable `ROERR_RAD1C`. Observation errors for 1C radiances are then written to the ODB in a call to **RADICOB** (from **HRETR_RAD**, which runs under the observation operator). However, these values are sometimes superseded by situation-dependent error schemes, such as those for the microwave sounders in **MW.CLEARSKY.OBERROR.MOD**. This is also done within the observation operator code.

For hyperspectral infrared sounder data, **DEFRUN** first specifies the default observation errors, and then overrides these by reading in ASCII files `rmtberr_airs`, `rmtberr_iasi` and `rmtberr_cris`. These files

specify observation errors separately for each channel, and correlations between them. The observation error specifications for AIRS and IASI are based on background and analysis departure diagnostics as explained in [Bormann *et al.* \(2016\)](#). For CrIS and IASI, the error covariances are estimated from samples of clear sky analysis departures computed in units of radiance. These are then mapped dynamically to brightness temperature (BT) units depending on the observed values for each given scene. This is to account for the nonlinearity of the Planck function which is important in the mid-wave and short-wave spectral regions, as explained in [Burrows and C.Lupu \(2024\)](#).

(b) *All-sky microwave radiances*

The observation error is determined using a ‘symmetric’ observation error model which is driven by the observed and first guess equivalent brightness temperatures. This can only be computed once the observation operator has been run, so it is computed within the observation operator code under the screening flag `LSCREEN=.TRUE`. As this is done within the observation operator code, more details are given in [Chapter 3](#).

(c) *AMVs*

Situation-dependent observation errors are computed in `AMV_OBERR`. These are generated within the AMV observation operator when `LSCREEN=.TRUE`.

(d) *GPSRO*

Observation errors are computed in `GPSRO_OBERROR`. These are generated within the observation operator when `LSCREEN=.TRUE`.

2.5.3 Ground-based radar precipitation composites

In the direct assimilation of NCEP Stage IV surface precipitation composites (using NEXRAD ground-based radars over the United States), observation error is expressed in $\log(\text{RR}[\text{mm}/\text{h}]+1)$ space and is assumed to have a constant value of 0.2 (resp. 0.1) for rainfall (resp. snowfall). This roughly corresponds to errors of 20% (resp. 10%) in terms of actual precipitation amounts. Practically, these errors are set in routines `GBRAD_PUT` and `GBRAD_PUT_TL`.

2.6 BACKGROUND ERROR ESTIMATES FOR SCREENING

All observations are assigned an estimate of the background error in observation space for later use in the background quality control (see [4.4.3](#)), and this estimate is stored in the ODB under `fg_error`. This estimate is only used to determine the expected variance of the background departures in the quality control against the background, and it is technically separate from the background error used during the assimilation for the control variables to determine the weighting of observations.

The assignment of the `fg_error` is performed in the routine `GEFGER`, and the method applied depends on the variable number `varno`. For the majority of geophysical variables the estimate is based on fields of situation-dependent estimates of the background error, calculated from the spread of the EDA. The fields are hence consistent with the derivation of the situation-dependent estimate of the background error used in the assimilation. The fields are available in MARS (scaled ensemble spread, SES), and they are read into variable `RZEGRID` in the routine `INIFGER`. In `GEFGER`, these fields are horizontally and - if required - vertically interpolated to observation locations. This method is applied for all standard geophysical variables (ie, T, q, TCWV, etc), and hence applies, for instance, to conventional observations and AMVs. Scaling of the error fields by `REDNMC` used in the assimilation is included. Note, however, that the SES fields are strictly only applicable for the initial time of the assimilation window, and no attempt is made to account for the temporal evolution of the background error pattern.

For observations with more complex observation operators, modifications of the above approach are used, dependent on whether estimates of the EDA spread have previously been calculated for the observed quantity. These result in estimates that may or may not be closely related to values of the actual statistical background error. This is, for instance, the case for:

Clear-sky radiances: For some clear-sky radiances, background error estimates based on the EDA spread (for a given zenith angle) are available as SES fields, and these are treated similar to the values for standard geophysical variables. SES fields are available for clear-sky simulations of HIRS, MSU, SSU, etc. (Bormann and Bonavita, 2013). Where needed, these values are interpolated to the observation locations. For some other sensors (e.g., geostationary imagers, ATMS), background error estimates from equivalent channels of the previously named sensors are assigned. While these estimates are not the same as a mapping of the actually used control vector background error into observation space they will capture similar spatial structures.

For hyperspectral infrared instruments (IASI, AIRS, CrIS), the background error estimate is set to 3 K. For most spectral regions this is a severe over-estimation of the uncertainty in the background, and the implication is that the background error quality control for these instruments is very loose.

All-sky radiances: All-sky radiances do not use situation-dependent estimates from the EDA spread, but rather employ a parametric model, as defined in [MWAVE.ERROR.MODEL](#).

Bending angles: The background error is set to the same value as the observation error.

2.7 OTHER PREPARATIONS

2.7.1 Surface emissivity for radiances

(a) Clear-sky microwave

Microwave ([EMIS_MW_N](#)) surface emissivities are set during the screening phase in [RAD1CEMIS](#) (called from [HRETR_RAD](#)) and stored in the ODB for later use by RTTOV. Setting the emissivity to values outside the range of 0 and 1 prompts the calculation of surface emissivity within RTTOV, using SURFEM-Ocean (Kilic *et al.*, 2023) for the microwave. This is done for all microwave radiances over sea.

For microwave radiances over land, several options exist to specify the surface emission, following the methods described in Karbou *et al.* (2006). The surface emissivity can be specified through an atlas, or it can be dynamically retrieved from window channel observations and FG estimates of skin temperature and atmospheric profiles, or the skin temperature can be retrieved given an emissivity atlas value and FG estimates of atmospheric profiles. Default choices are made by sensor in [SUEMIS_CONF](#) (including which channel is used for the last two options) and can be overwritten through the namelist [NAMEMIS_CONF](#). For the two options with dynamic retrieval of emissivity or skin temperature, the required radiative transfer calculations are performed in the routine [satrat/rttov/ifs/rttov_ec](#) when called from [HRETR_RAD](#) via [RADTR](#) or [RADTR_ML](#) (see also next section). The atlas values and the retrieved emissivities or skin temperatures are written to the ODB, and used as fixed input in subsequent calls to RTTOV. If atlas values are required, these are read in the routine [DEFRUN](#).

For some microwave sounders such as ATMS, a Kalman Filter is used to produce an evolving emissivity atlas from past dynamically retrieved emissivity values, as summarised in Krzeminski *et al.* (2009a) and Bormann (2014). The atlas is updated using the programs [EMISKF_UPDATE](#) in the [satrad](#) project (together with [emiskf*](#) and [kfgrid*](#) routines that can be found in the [emiss](#) directory of the [satrad](#) project). The program accesses the ODB and reads the required retrieved emissivity values. Only emissivity values that have the datum status flag “use_emiskf_only” set during the blocklisting are considered. Routine [EMISKF_INIT](#) specifies the resolution of the atlas and other configuration settings. Routine [EMISKF_INIT_ATLAS](#) is used to read the atlas values, [EMISKF_TRAJ](#) to evaluate the new emissivity values against the atlas values, [EMISKF_PREDICT](#) to predict forward in time the evolution of the errors in the atlas emissivity, and [EMISKF_UPDATE_ATLAS](#) to perform the update of the emissivity parametrization using the Kalman Filter equations, and [EMISKF_WRITE_ATLAS](#) to output the updated atlas. To use the atlas with a new sensor/channel, include the new sensor/channel in in the look-up table of known emissivity channels in [EMISKF_INIT](#), and provide a new ODBsql view in [EMISKF_UPDATE](#). The cycling of the atlas information is done through the files [emiskf.cycle*](#) which are stored as tar-ball in ECFS. The routine [EMISKF_INIT_ATLAS](#) is also used to read the atlas under [DEFRUN](#) during the application stage in the screening run. If no atlas is available a “coldstart” is performed, setting the atlas values and their errors to pre-specified values.

(b) *All-sky microwave*

A similar framework exists for all-sky microwave sounding radiances, but with the use of a fixed emissivity atlas (TELSEM) instead of an evolving Kalman Filter atlas. A dynamic emissivity retrieval is performed using representative frequencies with good sensitivity to the surface, taking into account the presence of light cloud and precipitation (Baordo and Geer, 2016). If the retrieval fails, which generally happens if the surface is not fully visible due to heavy cloud or precipitation, then a value is taken from the emissivity atlas. This process happens in the observation operator code under `MWAVE_EMIS`, when `LSCREEN=.TRUE.`

A limited number of all-sky microwave imager surface-sensitive channels are assimilated over land using a similar technique, but with the variation of surface emissivity as a function of frequency prescribed according to atlas. Another difference compared to the clear-sky microwave, applied to all all-sky sensors, is a more refined treatment of mixed scenes. For example, the scene is divided into water and land parts, and the surface emissivity water part is prescribed using the SURFEM-Ocean surface emissivity model, and the dynamic emissivity retrieval is carried out only in the part of the scene that is land. For further details see Geer *et al.* (2022).

Since cycle 49r1, some microwave imagers are now assimilated over sea ice surfaces (Geer, 2024a). Here, neither dynamic emissivity retrievals nor atlases are used, but rather there is an empirical sea ice surface emissivity model (Geer, 2024b). The input parameters of the empirical surface emissivity model, which are a statistical representation of the snow and ice state, are retrieved as part of 4D-Var data assimilation, at observation locations, along with the sea ice concentration.

(c) *Infrared*

Similarly to the microwave, infrared (`EMIS_IR`) surface emissivities are set during the screening phase in `RADICEMIS`. For infrared IASI radiances over land, the surface emissivity is specified through an atlas, provided by the University of Wisconsin infrared land surface emissivity atlas (UWIREMIS). For other infrared instruments, a single constant emissivity value of 0.98 over land is used for all wavelengths. For ocean, ISEM-6 (Sherlock, 1999) is used, including sea-ice areas.

2.7.2 AMV height reassignment

(a) *Legacy height reassignment scheme for diagnostics*

A framework exists to perform height reassignment for AMVs when `LSCREEN=.TRUE.`. This is not used operationally. The reassignment is done using the routine `AMV_REASSIGN`, called from `HRETR_CONV`.

(b) *Operational reassignment of low-level AMVs*

In order to mitigate height assignment problems of AMVs in cases of high wind shear, some AMVs have their heights reassigned to the model cloud layer average. The scientific motivation for the height reassignment scheme, and evidence of its beneficial impact on forecast skill, are explained in Lean and Bormann (2021). The height reassignment scheme is used via the `LAMV_HEIGHT_ADJUST=.TRUE.` switch.

Model cloud tops and cloud bases are provided by the routine `AMV_GET_MODCLOUD`, called from `HRETR_CONV`. A model level is considered cloudy if its model cloud fraction is above 1%, and the sum of its cloud liquid and ice water content is greater than 10^{-6} kg/kg. Three consecutive levels must meet the cloudy criteria for the routine to record a cloud top and cloud base and therefore for the height reassignment to be applied to a given AMV.

The reassignment of AMV heights is performed in the routine `HRETR_CONV`. Within this routine further conditions are checked before the reassignment scheme is applied: the AMV must be between 700 and 900 hPa, and must be above the model cloud top. Finally, the scheme is only applied for dates later than 31 December 2009, this date cut-off ensures that the height reassignment scheme is only operating when low level clouds are sufficiently represented in the model.

The pre-reassignment AMV height is copied to the variable `VERTCO.REFERENCE_2` while the height stored in `VERTCO.REFERENCE_1` is modified. The model u and v wind values at the pre-reassignment height are stored in ODB columns `UMOD_OLD` and `VMOD_OLD`; this is so that O-B statistics can be derived at either height. Optionally, the model cloud-top and cloud-base pressure can be written to the ODB in the `CT_P` and `CB_P` columns.

Chapter 3

Observation operators

Table of contents

3.1 Introduction

- 3.1.1 Top-level observation operator: HOP
- 3.1.2 Tangent-linear and adjoint code
- 3.1.3 Data selection controls: NOTVAR
- 3.1.4 Test harness
- 3.1.5 Other tests of the observation operator

3.2 Conventional observation operators

- 3.2.1 General aspects
- 3.2.2 Geopotential height
- 3.2.3 Wind
- 3.2.4 Humidity
- 3.2.5 Temperature
- 3.2.6 Surface observation operators
- 3.2.7 Atmospheric Motion Vectors
- 3.2.8 Gas retrievals
- 3.2.9 Scatterometer winds

3.3 Satellite radiance operators

- 3.3.1 Common aspects for the setup of nadir radiance assimilation
- 3.3.2 Clear-sky nadir radiances
- 3.3.3 All-sky nadir radiances
- 3.3.4 Clear-sky limb radiances

3.4 GPS Radio Occultation bending angles

3.5 Ground-based radar precipitation composites

3.6 Atmospheric composition

3.7 Cloud radars and lidars

- 3.7.1 Observation operator
- 3.7.2 Observation error, screening and bias correction

3.1 INTRODUCTION

The observation operators provide the link between the analysis variables and the observations ([Lorenc, 1986](#); [Pailleux, 1990](#)). The observation operator is applied to components of the model state to obtain the model equivalent of the observation, so that the model and observation can be compared. The operator H signifies the ensemble of operators transforming the control variable x into the equivalent of each observed quantity, y^o , at observation locations.

From a data assimilation perspective, the job of the observation operator is to interpolate the model state to the observation location, and then to convert from state variables to observed variables. In the IFS, this job is divided into two or three parts. The horizontal interpolation to observation locations is handled by the GOMs; this is a complex task requiring multi-process communication to interpolate and to get the model data from whichever processors it is located onto the processors where the observation data resides (see part 2). From a dataflow point of view, the observation operator is needed in the context




Figure 3.1 Simplified IFS flow diagram for observation-related processing in 4D-Var. This is an illustrative, rather than comprehensive, description of the processing. See Fig. 1.1 for the wider observation processing context.

of observation preparation, screening and 4D-Var. Figure 3.1 summarises the operations relevant to observations in 4D-Var. In this observation operator documentation, we consider a more limited part of the observation operator: the vertical part of the interpolation (where relevant) and the transformation into observed variables.

The IFS observation operators are generic in the sense that the same routine is often used for several different data types. For example, the radiance operator (RTTOV) simulates measurements from a large number of satellite radiometers (microwave and infrared), and the temperature operator (PPT) is used for TEMP, AIREP, and other data types. Similarly the routine PPQ is used for interpolation of specific humidity to given pressures, but it can also be used for any other atmospheric mixing ratio constituents, such as ozone and carbon dioxide. Note that many of the PP-routines were developed for the model’s pressure-level post-processing package and are used also in that context.

3.1.1 Top-level observation operator: HOP

HOP runs the observation operator for one set (see Sec. 1.2.2) of observations. The inputs are:

- the GOM.PLUS (YDGP5) containing the model state at observation locations;
- the ODB, currently an implicit input through a Fortran module; as OOPS develops further the ODB will become an explicit input too;
- the VarBC object YDVARBC, representing the bias correction coefficients;
- the implicit (communicated-by-module) results of a variety of setup routines (see later).

The outputs are:

- the observation equivalent $H(\mathbf{x})$;
- the bias correction in observation space;
- the ODB is updated (mainly only when LSCREEN=.TRUE.) to store the results of operator-level screening, prior retrievals (done in HRETR routines, e.g. surface emissivity) and operator-level observation error assignments.

The main job of HOP is first to open an ODB view, i.e. to give access to the contents of the ODB relevant to this set of observations, and second, to call a lower level observation operator depending on what observational data is in the set. This is determined by the varnos (see Sec. 6.3) of the reports and data in the set. The varno determines what physical observation operator is required, for example radiance simulation using RTTOV or a vertical interpolation of humidity. Because observation sets often contain multiple variables, it is possible for more than one physical observation operator to be called for the same set. The main physical operators are:

- **Conventional observations (OBSOP_CONV)**: this covers surface and upper-air observations, as well as in-situ data and satellite retrievals that include AMVs and scatterometer winds. What links these observations is the observed variables are closely related to the model state, so that the main part of the observation operator is usually only a vertical interpolation or integration (for example interpolating to the vertical level of an aircraft observation, or integrating over many model levels for total or partial column ozone observations). All the necessary conventional operators are handled by 'PP' or post-processing routines, and are accessed through the top-level PP interface to the observation operator, **PPNEW**. Section 3.2 covers these observation operators.
- **Nadir radiances (OBSOP_RAD)**: this covers both clear-sky and all-sky radiance observations. All radiances uses RTTOV. Scattering for MW sensors now run through the main RTTOV interface in a very similar way to IR sensors (a separate RTTOV-SCATT model no longer exists). See section 3.3
- **Limb radiances (OBSOP_LIMB)**: the 2-dimensional model state and 2-dimensional observation operator for limb radiances is not handled by RTTOV, so limb radiances need their own operator. This code is dormant and not used operationally. See section 3.3
- **GPS bending angles (OBSOP_GPSRO)**: this covers satellite radio-occultation, and it usually requires 2-D slice through the model to compute the tangent path through the earth's atmosphere, including the computation of the bending angles. See section 3.4
- **GPS path delays (ground-based GPS; OBSOP_GPS_SURFACE)**: Currently only active at Météo-France, so not discussed further.
- **Ground-based radar observations (OBSOP_RADAR)**: This implements the 1D-Bayesian retrieval of [Wattrelot et al. \(2014\)](#) used at Météo-France, and not discussed further here.
- **Precipitation accumulations, (OBSOP_PRECIP_ACCUM)**: This enables the assimilation of NEXRAD rain radar and rain gauges. The model-equivalent precipitation accumulations are computed inside the model physics and then stored directly to ODB, circumventing the normal process of the observation operator. The main job of the observation operator under HOP is to extract the pre-computed observation equivalent from the ODB.
- **Atmospheric composition (OBSOP_COMPOSITION)**: This handles observations mainly used in the Copernicus atmospheric monitoring service, and it covers retrieved constituent variables as well as aerosol optical depth and some other advanced experimental operators. Composition variables would normally be expected to go through the conventional observation route (OBSOP_CONV). However, the composition route allows proper treatment of the averaging kernels. A more technical link among the composition operators is their use of the 'GEMS' input model variables, split into greenhouse gases (GHG), reactive gases (CHEM) and aerosols (AERO). Nevertheless, there is much overlap at a scientific and technical level between the conventional and composition parts of the observation operator, so a long-term design goal should be to unify them.
- **Cloud radar and lidar (OBSOP_RADLID)**: This path supports EarthCARE radar and lidar as well as earlier Cloudsat and Calipso radar and lidar missions, and simulates radar reflectivities and lidar backscatter.

3.1.2 Tangent-linear and adjoint code

In the upper levels of the observation operator, tangent-linear (TL) and adjoint (AD) routines are often merged with the direct (i.e. nonlinear) version. The top-level observation operator routine, **HOP** is one example. Within each of these routines the mutually exclusive logicals LLDIRECT, LLTL and LLAD indicate the mode of operation. TL or adjoint mode is activated by the presence of optional arguments containing the additional data structures needed in a TL or adjoint context. Whether present or not, these optional arguments can cascade down the calling tree. For example, the GOM.PLUS.TL structure

(model increment at observation locations) is an optional argument to HOP, which is then optionally passed down to the lower level operators.

In the lower levels of the code, and particularly those that do extensive numerical computation, the direct, TL and adjoint will usually be three separate routines. Merging the three options together into one routine makes sense in areas where little numerical computation is done: in many of the routines that have been merged, 90% of code was common across the direct, TL and adjoint. Merging the routines reduces the maintenance overhead (preventing changes having to be implemented three times over) and makes it easier to keep the direct, TL and adjoint routines consistent. For example, if a new feature is added in the direct version, it is more obvious that the TL and adjoint should be updated to reflect this.

Where direct and TL or adjoint variables are present in the same routine, the usual IFS naming convention applies:

- VARIABLE5 is the nonlinear (direct) variable;
- VARIABLE is the TL or adjoint variable, with its meaning depending on the context (e.g. LLTL or LLAD);

3.1.3 Data selection controls: NOTVAR

Most data selection criteria are kept in the blacklist or are performed in observation operator code under the LSCREEN=.TRUE. switch. (see later). However, classes of data can also be switched on and off in the observation operator using the NOTVAR array in **NAMJO**. The second dimension in this array is the observation type. The first dimension is the variable number (varno, see later). The elements of the NOTVAR array can take either of two values: 0, means that the data will be used; -1, means that it won't. NOTVAR is only intended for operational emergencies or debugging, as uniquely it can stop some observations going through the observation operators. The blacklist only stops observations going into the 4D-Var minimisation. If an observation is crashing the observation operator in the screening trajectory, NOTVAR is a temporary solution; later the code should be re-written to prevent the crash, or the data should be blacklisted in the usual way.

3.1.4 Test harness

The test harness for HOP, **HOP_DRIVER**, is a program that can be run independently of the IFS with suitable input files:

- GOM.PLUS.DUMP files containing the model state at observation locations;
- ODB-1 files (ECMA and CCMA) containing the ODB;
- VARBC.cycle file containing the VarBC coefficients;
- HOP_RESULTS files containing reference HOP output, for comparison;
- All the data files required to run the observation operator, such as RTTOV coefficient files, channel mappings etc.

To produce these files it is necessary to run the IFS with LHOP_RESULTS=true and LGOM.PLUS.DUMP=true in the TESTVAR namelist. Also it is necessary to restrict the amount of observational data to make it feasible to run the test harness interactively on a single node. Further details can be found in the test harness **README.run** file.

In direct mode the test harness runs the observation operator and computes the VarBC bias correction; these are compared to the reference output and should be bit-identical between offline and online runs of the same code. In TL and adjoint mode the TL-adjoint consistency test is run, along with a comparison of the TL results with reference results from a run of the IFS. There is not yet the facility to test VarBC TL or adjoint or the auxiliary control variables like the skin-temperature sink variable.

3.1.5 Other tests of the observation operator

The adjoint test of the whole 4D-Var operator chain (change of variable, forecast model, interpolation, observation operator) can be activated with NTESTVAR = 1 and LADTEST = .TRUE. There are actually

two tests of adjoint triggered here; the relevant one for observations is coded under `ADJOTEST`. The results of this test are printed to the logfiles after ‘Adjoint test obs. operator 1.234567890123456789’.

An adjoint test of the RTTOV clear-sky observation operator, on a per-set basis, can be activated with `NTESTVAR = 1` and `LRTTOV_ADTEST = .TRUE`.

3.2 CONVENTIONAL OBSERVATION OPERATORS

3.2.1 General aspects

A review of the processing/assimilation of in situ observations is provided by [Pauley and Ingleby \(2022\)](#). Two of the “OBSTYPE” names in the IFS are based on old code formats and are misleading: the AIREP grouping includes all aircraft data and the DRIBU group would be better called BUOY. (Ship data is included in the SYNOP group along with METARs). For in situ observations the station identifier (STATID) is important - used in duplicate checks, blocklisting, bias correction and monitoring. SYNOPs and radiosondes have traditionally used a five-digit identifier, but there is a move by WMO to more complex/longer WIGOS identifiers (WIGOSID). Complex WIGOSIDs are mapped to unique 8-character identifiers (stored in STATID) by the ECMWF SAPP system; in this way bias correction etc continues to use STATID, but the externally recognised WIGOS identifier is also stored. Reports can be received either in Traditional Alphanumeric Codes (TAC) or the newer binary BUFR templates. Duplicate checks prefer the BUFR reports (unless they are blocklisted), for radiosondes the report with most levels is preferred.

For radiosondes BUFR reporting provides the time, latitude and longitude displacement of each point relative to the start/launch point. TAC profiles are treated as vertical and instantaneous, but since 2018 BUFR profiles have been split into sub-profiles (initially 15-minutes long for ascents, 5-minutes for dropsondes and descent profiles, now 5-minutes for ascents too). Each sub-profile is treated as being valid at the latitude/longitude/time of its lowest point. (The sub-profiles are identified by SEQNO values - whole reports by REPORTNO values. Select all records with the same REPORTNO value to reconstruct a report. Using the archive select on analysis date/time too.) This treatment of radiosonde drift can improve the fit to background values by 5-10% in the stratosphere ([Pauley and Ingleby, 2022](#)). Vertical thinning of radiosondes uses a factor ZFAC, default 3. The first level with pressure p_0 in Pa is chosen, we then search for the next level with $|p_1 - p_0| > 100 * ZFAC$ or $|p_1 - p_0| / 0.5(p_1 + p_0) > 0.01 * ZFAC$ and so on. There is an override, so that standard levels are always chosen. For a full profile this gives about 350 levels, approximately 3 per model layer. From 2019 BUFR DROPSONDE data were assimilated, this necessitated a tailored duplicate check (based on start positions and times). Initially profiles were considered duplicates if the times were within 10 minutes and latitude/longitude within 0.1 degree (specified in `RDUPT_DROP` and `RDUPD_DROP`). Following a case in summer 2024 with excessive density of dropsonde data assimilated these values were changed to 90 minutes and 0.6 degrees, and ZFAC for dropsondes was set to 9.

This section describes the great variety of observation operators handled by `OBSOP_CONV` and underneath, `PPNEW`, which is a wrapper for most of the PP routines. Some details can be found in [Courtier et al. \(1998\)](#), and in the following sub-sections. The vertical operations depend on the variable. The vertical interpolation is linear in pressure for temperature (`PPT`) and specific humidity (`PPQ`), and it is linear in the logarithm of pressure for wind (`PPUV`). The vertical interpolation of geopotential (`PPGEOP`) is similar to wind (in order to preserve geostrophy) and is performed in terms of departures from the ICAO standard atmosphere for increased accuracy ([Simmons and Chen, 1991](#)). The current geopotential vertical interpolation together with the temperature vertical interpolation are not exactly consistent with the hydrostatic equation. A new consistent and accurate vertical interpolation has been devised by Météo-France, which may be important for intensive use of temperature information. The new routines have been tested by ECMWF and as the results were not unambiguously positive the new routines have not yet been adopted – and they are not described in this documentation. In the meantime, the old routines are still used (switch `LOLDPP = .TRUE` in `namct0`), under the names `PPT_OLD`, `PPGEOP_OLD` and `PPUV_OLD`, with tangent linear `PPTTL_OLD`, `PPGEOPTL_OLD` and `PPUVTL_OLD` and adjoint `PPTAD_OLD`, `PPGEOPAD_OLD` and `PPUVAD_OLD`.

The vertical gradients of the model variables vary strongly in the lowest part of the boundary layer, the vertical interpolation operator for those data takes this into account following Monin–Obukhov

similarity theory [Geleyn \(1988\)](#). The vertical interpolation operators for SYNOP 10-metre wind, 2-metre temperature and 2-metre specific humidity are documented in Chapter 3.9 of Part 4 of the IFS documentation. Values are stored in the GOM arrays and accessed directly by `OBSOP_CONV`. Various changes at 49r1, including the introduction of 2-metre temperature assimilation in 4D-Var are described in [Ingleby et al. \(2024\)](#). (For now the 2-metre values are only assimilated for the first six hours of lond window assimilation but this restriction may be relaxed in future. Prior to 49r1, and the use of `L_SCREEN_LEVEL_OBSOP=.TRUE.`, an older version of the boundary layer interpolations was used see below.) The Monin–Obukhov based observation operator for 10-metre wind is used for all surface winds (SYNOP, BUOY, TEMP, PILOT and SCAT). For BUOY data wind interpolation is performed to the actual observation height (in practice ranging from 4 to 10 m). Ship winds higher than 10 m are assimilated by interpolation of the model layers above and below the level.

The model equivalent of total column water vapour (TCWV) is obtained by vertical integration of q (in `GPPWC` and `PPPWC`). The routine `PPPWC` is also used for vertical integration of GEMS/MACC trace gasses. Observation operators exist for precipitable water content (also using `PPPWC`) and thicknesses (`PPGEOP`). At 49r1 we are starting the assimilation of zenith total delay (ZTD) from ground-based GNSS stations (known in the IFS as GroundGPS, the VARNO is "atmospheric path delay in satellite signal" or APDSS), providing information on TCWV. The data are bias corrected using VarBC with three predictors: constant, 1000-300 hPa thickness and model TCWV (based on Météo-France usage, they have used the data for almost 20 years). The STATID is made up of two parts eg 'ABANIGE2' - the station is ABAN and the processing is by IGE2. We can receive data from the same station that has been through several different processing systems. A long list of allowed STATIDs is specified in the data selection and updated occasionally, this is to try to pick the best processing stream at each station and also introduce some thinning where the stations are very dense.

3.2.2 Geopotential height

The geopotential at a given pressure p is computed by integrating the hydrostatic equation analytically using the ICAO temperature profile and vertically interpolating $\Delta\phi$, the difference between the model level geopotential and the ICAO geopotential ([Simmons and Chen, 1991](#)). The ICAO temperature profile is defined as

$$T_{\text{ICAO}} = T_0 - \frac{\Lambda}{g} \phi_{\text{ICAO}} \quad (3.1)$$

where T_0 is 288 K, ϕ_{ICAO} is the geopotential above 1013.25 hPa and Λ is 0.0065 K m⁻¹ in the ICAO troposphere and 0 in the ICAO stratosphere (the routine `PPSTA`). The ICAO tropopause is defined by the level where the ICAO temperature has reached 216.5 K (`SUSTA`). Using this temperature profile and integrating the hydrostatic equation provides T_{ICAO} and the geopotential ϕ_{ICAO} as a function of pressure (`PPSTA`). We may then evaluate the geopotential $\phi(p)$ at any pressure p following

$$\phi(p) - \phi_{\text{surf}} = \phi_{\text{ICAO}}(p) - \phi_{\text{ICAO}}(p_{\text{surf}}) + \Delta\phi \quad (3.2)$$

where p_{surf} is the model surface pressure and ϕ_{surf} , the model orography. $\Delta\phi$ is obtained by vertical interpolation from the full model level values $\Delta\phi_k$. The interpolation is linear in $\ln(p)$ up to the second model level (`PPINTP`) and quadratic in $\ln(p)$ for levels above it (`PPITPQ`, see below). Following [Simmons and Burridge \(1981\)](#) the full model level values are obtained by integrating the discretized hydrostatic equation using the routine `GPGEO` of the forecast model to give

$$\Delta\phi_k = \sum_{j=L}^{k+1} R_{\text{dry}}(T_{v_j} - T_{\text{ICAO}_j}) \ln\left(\frac{p_{j+1/2}}{p_{j-1/2}}\right) + \alpha_k R_{\text{dry}}(T_{v_k} - T_{\text{ICAO}_k}) \quad (3.3)$$

with

$$\alpha_k = 1 - \frac{p_{k-1/2}}{p_{k+1/2} - p_{k-1/2}} \ln\left(\frac{p_{k+1/2}}{p_{k-1/2}}\right)$$

for $k > 1$ and $\alpha_1 = \ln(2)$.

(a) *Quadratic vertical interpolation near the top of the model*

Above the second full level of the model, the linear interpolation (PPINTP) is replaced by a quadratic interpolation in $\ln p$, performed in the routine PPITPQ using

$$z(\ln p) = a + b(\ln p) + c(\ln p)^2 \quad (3.4)$$

where a , b and c are constants determined so that the above equation fits the heights at the top levels ($k = 1, 2$ and 3). The interpolation formula is

$$\phi(\ln p) = z_2 + \frac{(z_2 - z_1)(\ln p - \ln p_2)(\ln p - \ln p_3)}{(\ln p_2 - \ln p_1)(\ln p_1 - \ln p_3)} - \frac{(z_2 - z_3)(\ln p - \ln p_1)(\ln p - \ln p_2)}{(\ln p_2 - \ln p_3)(\ln p_1 - \ln p_3)} \quad (3.5)$$

where 1, 2 and 3 refer to levels $k = 1, 2$ and 3 , respectively.

(b) *Below the model's orography*

The extrapolation of the geopotential below the model's orography is carried out as follows: Find T^* (surface temperature) by assuming a constant lapse rate Λ , from the model level above the lowest model level (subscript $l - 1$), see the routine CTSTAR, using

$$T^* = T_{l-1} + \Lambda \frac{R_{\text{dry}}}{g} T_{l-1} \ln \frac{p_{\text{surf}}}{p_{l-1}} \quad (3.6)$$

$$T^* = \frac{\{T^* + \max[T_y, \min(T_x, T^*)]\}}{2} \quad (3.7)$$

Find the temperature at mean sea level, T_0 (also in CTSTAR) from

$$T_0 = T^* + \Lambda \frac{\phi_{\text{surf}}}{g} \quad (3.8)$$

$$T_0 = \min[T_0, \max(T_x, T^*)] \quad (3.9)$$

where T_x is 290.5 K and T_y is 255 K. The geopotential under the model's orography is (in PPGEOP) calculated as

$$\phi = \phi_{\text{surf}} - \frac{R_{\text{dry}} T^*}{\gamma} \left[\left(\frac{p}{p_{\text{surf}}} \right)^\gamma - 1 \right] \quad (3.10)$$

where $\gamma = \frac{R_{\text{dry}}}{\phi_{\text{surf}}} (T_0 - T_{\text{surf}})$.

3.2.3 Wind

In PPUV a linear interpolation in $\ln p$ (PPINTP) is used to interpolate u and v to the observed pressure levels up to the second full model level, above which a quadratic interpolation is used (PPITPQ). Below the lowest model level wind components are assumed to be constant and equal to the values of the lowest model level.

3.2.4 Humidity

Specific humidity q , relative humidity U and precipitable water content PWC are linearly interpolated in p , in PPQ, PPRH and PPPWC, respectively. Upper air relative humidity data are normally not used, but could be used, if required. The use of surface relative humidity data is described in Subsection (a).

(a) *Saturation vapour pressure*

Radiosonde and surface reports usually contain humidity as dew point temperature (some surface reports contain relative humidity instead). The use of dew point i.e. water-phase saturation is in accordance with the WMO rules. For radiosondes WMO (2024) (Chapter 12) recommends the use of several similar saturation vapour pressure equations, we choose to use that from Sonntag (1994). This is implemented in statement function FOEESONNTAG (in function/stmfun_ifs.func.h), the forecast

model uses FOEEWM (different phases). The observed humidity is converted to specific humidity for assimilation (also relative humidity for diagnostic purposes). The use of the Sonntag equation was introduced in 2016 and gave improved humidity fit to the background in the upper troposphere for the best radiosonde types (Ingleby, 2017).

(b) *Relative humidity*

In **GPRH** relative humidity U is computed from

$$U = \frac{pq \frac{R_{\text{vap}}}{R_{\text{dry}}}}{\left[1 + \left(\frac{R_{\text{vap}}}{R_{\text{dry}}} - 1\right)q\right] e_{\text{sat}}(T)} \quad (3.11)$$

and then in **PPRH** interpolated to the required observed pressure levels (using **PPINTP**). Below the lowest model level and above the top of the model is U assumed to be constant. Saturation vapour pressure is calculated using FOEEWMO if **GPRH** has been called from the observation operator routines, and using FOEEWM if called from the model post processing.

(c) *Precipitable water*

In **GPPWC** precipitable water is calculated as a vertical summation from the top of the model by

$$PWC_k = \frac{1}{g} \sum_{i=1}^k q_i (p_i - p_{i-1}) \quad (3.12)$$

and then in **PPPWC** interpolated to the required observed pressure levels (using **PPINTP**). PWC is assumed to be zero above the top of the model. Below the model's orography PWC is extrapolated assuming a constant $q = q_l$.

(d) *Specific humidity*

Specific humidity q is in **PPQ** interpolated to the required observed pressure levels (using **PPINTP**). Below the lowest model level and above the top of the model is q assumed to be constant and equal to q_l and q_1 , respectively.

3.2.5 Temperature

Temperature is interpolated linearly in pressure (**PPINTP**), in the routine **PPT**. Above the highest model level the temperature is kept constant and equal to the value of the highest model level. Between the lowest model level and the model's surface the temperature is interpolated linearly, using

$$T = \frac{(p_{\text{surf}} - p)T_l + (p - p_l)T^*}{p_{\text{surf}} - p_l} \quad (3.13)$$

Below the lowest model level the temperature is extrapolated by

$$T = T^* \left[1 + \alpha \ln \frac{p}{p_{\text{surf}}} + \frac{1}{2} \left(\alpha \ln \frac{p}{p_{\text{surf}}} \right)^2 + \frac{1}{6} \left(\alpha \ln \frac{p}{p_{\text{surf}}} \right)^3 \right] \quad (3.14)$$

with $\alpha = \Lambda R_{\text{dry}}/g$, for $\phi_{\text{surf}}/g < 2000$ m, but α is modified for high orography to $\alpha = R_{\text{dry}}(T'_0 - T^*)/\phi_{\text{surf}}$, where

$$T'_0 = \min(T_0, 298) \quad (3.15)$$

for $\phi_{\text{surf}}/g > 2500$ m, and

$$T'_0 = 0.002[(2500 - \phi_{\text{surf}}/g)T_0 + (\phi_{\text{surf}}/g - 2000) \min(T_0, 298)] \quad (3.16)$$

for $2000 < \phi_{\text{surf}}/g < 2500$ m. If $T'_0 < T^*$ then α is reset to zero. The two temperatures T^* and T_0 are computed using (3.6) to (3.9).

3.2.6 Surface observation operators

Preparations for the vertical interpolation of surface data are done in during the creation of the GOM_PLUS. Here dry static energy (**SURBOUND**), Richardson number, drag coefficients and stability functions (**EXCHCO**) are computed. For scatterometer data, information on equivalent neutral 10-metre wind is directly fetched from the physics package (via the GOM arrays) which is preprocessed in routine **EXCHCO_VDF**. The actual vertical interpolation is performed in **PPOBSAS**, which embraces routines for 10-metre vector-wind components (**PPUV10M**), 2-metre temperature (**PPT2M**) and 2-metre relative humidity (**PPRH2M**).

The routines mentioned above for 2-metre temperature and humidity and 10-m wind were superceded for postprocessing some years ago by **SPPCFL_MOD**, documented in Chapter 3.9 of Part 4 of the IFS documentation. At CY49R1 the newer routines became the default in the surface observation operators for SYNOP data (**L_SCREEN_LEVEL_OBSOP = .TRUE.**). The 2-metre temperature (or humidity) is interpolated between the value at the lowest model level and the surface value. Note that 2-metre relative humidity (which still uses the old observation operator) is replaced by 2-metre specific humidity as the assimilated variable. There is a technical issue in that model equivalents for the first time step (15 minutes) are not available so for now these values are blocklisted if the new option is used. The namelist variable (**4DLAPSET2**, default 0.0055 K/m) is the lapse rate used in 4D-Var for T2m to adjust for the difference between station height and model orography (if **L_SCREEN_LEVEL_OBSOP** is true). Similar adjustments (with a lapse rate of 0.0065 K/m) are used in other NWP systems (eg [Ingleby \(2015\)](#)) and also in the verification against SYNOPs at ECMWF. The tangent linear and adjoint computations are still based on the old scheme from ([Geleyn, 1988](#)). Currently the outer loop interpolation for BUOY data still uses the old scheme. One minor advantage of the old scheme is that it allows for the direct interpolation of winds to anemometer heights below 10 m, whereas the new scheme is currently hardwired to 10 m.

(a) Two-metre relative humidity

In **GPRH** relative humidity is computed according to (3.11). The relative humidity depends on specific humidity, temperature and pressure (q , T and p , respectively) at the lowest model level. It is constant in the surface model layer, see **PPRH2M**.

3.2.7 Atmospheric Motion Vectors

Groups of AMVs (aka SATOBs) are set up in the routine **SUAMV**, one group per satellite, computational method, and codetype. The information is stored in the satob group table, residing in the module **YOMSATS**.

The group table also specifies what type of observation operator is to be used for the particular group (entry `obs_oper`). The default used in operations is to assimilate all AMVs as single-level wind observations ([Tomassini et al., 1998](#); [Bormann et al., 2003](#)), in the same way as conventional data, using the same interpolation routine. Other options are treating AMVs as layer averages, with weights defined, for instance, by a boxcar weighting function.

3.2.8 Gas retrievals

Retrievals of atmospheric species such as ozone or water vapour are used in the form of integrated layers bounded by a top and bottom pressure which are given as a part of the observation. The same observation operator is used as for precipitable water (**PPPWC**). The same concept is applied to all data, whether it is total column data (like TOMS and GOME ozone data or MERIS total column water vapour) or data with higher vertical resolution (like SBUV). For ozone, variational bias correction is implemented (see Part II, class name “to3”, module **VARBC_TO3**). SBUV data is currently used as anchor for the variational bias correction of ozone and therefore assimilated without bias correction.

3.2.9 Scatterometer winds

For ERS-2 and ASCAT normally two ambiguous pairs of u -component and v -component observations are found at each SCAT location – with directions approximately 180 degrees apart. QuikSCAT can

have 2, 3 or 4 ambiguous winds. Up to the first NSCAWSOLMAX (4 by default, adaptable through the namelist `NAMSCC`) wind solutions are accepted. In case only one ambiguity is found, the report is rejected. If `LQSCATT = .TRUE.` (the default, modifiable through the namelist `NAMJO`), the normal quadratic J_o will be used. In this case only the SCAT wind nearest to the high resolution background will be used (which is determined in a section of `HOP`). For winds that are not closest to the first guess or analysis, global datum event flag 9, respectively 10 is set (see Table 6.24). For the latter case datum status is set to rejected as well (Table 6.23). When `LQSCATT = .FALSE.`, the two first winds are used and the ambiguity removal takes place implicitly through a special SCAT cost-function (see part II), in `HJO` (Stoffelen and Anderson, 1997). In that case for QuikSCAT the most likely wind (highest *a priori* probability) and its most opposing ambiguity are selected.

Routine `PPUV10M` is used for SCAT data as well as for SYNOP, SHIP and BUOY wind. The difference is that (in case account is taken for ocean current) the relative wind rather than the absolute wind is returned (not active in the operational suite, though), and the evaluation of equivalent-neutral wind rather than the real wind (which latter includes for scatterometer data undesired stability effects; operational since November 2011).

3.3 SATELLITE RADIANCE OPERATORS

The majority of satellite data assimilated currently is radiances, all of it going through the top-level observation operator `OBSOP_RAD`. Radiances, rather than retrieved products, are assimilated directly (Andersson *et al.*, 1994), wherever possible. The current operational configuration uses clear level-1C radiances from a number of sensors, including ATOVS (McNally *et al.*, 1999), AIRS, IASI, CrIS, ATMS, as well as geostationary water vapour clear-sky radiances (Munro *et al.*, 2004).

Clear, cloudy and precipitation-affected radiances from microwave imagers, humidity sounders, and temperature sounders (e.g. AMSU-A) are assimilated using an all-sky approach. Scattering simulations for MW sensors are now done through the main RTTOV radiative transfer model (Saunders and Matricardi, 1998; Matricardi *et al.*, 2001; Saunders *et al.*, 2018), currently using RTTOV version 14.0. This is a significant update to RTTOV as all science capabilities from the old RTTOV-SCATT model are now implemented and enhanced within RTTOV. This enables all visible/IR/MW scattering simulations to be run through the RTTOV interface, and most RTTOV science options are now available across the spectrum. Other important features of RTTOV-14 are an improved representation of the atmospheric profile that better matches the way that NWP models represent the atmosphere and improvements related to the treatment of surface emissivity and reflectance, including the capability to represent heterogeneous surfaces (for example, land/sea or sea/sea-ice) within the satellite field of view.

Within the IFS there are two separate observation operator paths for nadir radiances, one for clear-sky radiances (including optionally all-sky infrared) and one for all-sky microwave radiances. These separate paths start within `OBSOP_RAD`. This section also covers the (non-operational) satellite limb radiances that are handled under `OBSOP_LIMB` as they have many commonalities.

3.3.1 Common aspects for the setup of nadir radiance assimilation

The operational radiance assimilation shares the following setup aspects for both radiance assimilation routes. The datasets are distinguished by a satellite ID, a sensor ID, and a codetype. The latter is used to distinguish clear-sky (codetype=210=NGTHR) or all-sky radiances (codetype=215=NSSMI).

The main set-up routine for radiances is `SURAD`. It recognises BUFR satellite IDs (call to `GETSATID`), reads RTTOV coefficient files (call to `RTSETUP`), and builds a “satellite group table” containing information on which satellite groups are present.

(a) Satellite identifiers and sensors

Satellite identifiers are dealt with in the routine `GETSATID`, called from `SURAD`. The ODB contains the identifiers as given in the original BUFR messages. Lists of identifiers for which data exist in any given ODB are prepared in the routine `SURAD`. The routine `GETSATID` matches those BUFR satellite identifiers with the platform and satellite numbers used by the RT-code (e.g. platform 1, satellite 10 for

NOAA-10). New satellite IDs need to be defined in `GETSATID`, but the id-conversion tables can also be modified through the namelist `NAMSATS`.

The various types of radiance data in the IFS are also classified by sensor. Each satellite sensor is assigned a number, defined in the module `YOMSATS`. The sensor number is used as index to various tables containing observation errors, BgQC thresholds, VarQC parameters, the J_o -table `JOT`, etc. See the routine `DEFRUN`.

(b) *Satellite group table*

Various satellite-related indices are gathered in the routine `SURAD` in the FORTRAN90 data structure called the 'satellite group table', `satgrp.t` (defined in `YOMSATS`). The table contains elements such as the satellite ID, the sensor ID, the codetype, a sequence number for addressing the transmittance coefficients (`rtcoef_pos`), the number of channels, a channel number list, etc. The various satellite-related indices are universally determined across all processors. There is one entry in the satellite group table per satellite, sensor, and codetype. A list of all the satellite groups that were found in the ODBs can be found in the `ifstraj` output by searching for `SATGRP`.

(c) *Radiative transfer coefficients, pressure levels and validation bounds*

Various preparations for RTTOV calculations are set-up in the call to `RTSETUP` from `SURAD`. This includes reading of the various coefficient files required by RTTOV. The files can be found under `/home/rd/rdx/data/ < cycle > /sat/rttov`. There is one file with the prefix `rtcoef_` containing coefficients for the general clear-sky radiative transfer for each instrument and satellite, and files with prefix `hydrotable_` used for a parametrization of cloud scattering effects for some infrared sensors, and for scattering lookup tables for all-sky microwave. Both types of files are read via a call to `RTTVI` (in the `satrad` library) which also communicates other settings from the RTTOV code in the `satrad` library to the IFS side. Only the files that are required are read, i.e. only the files for which observations are present (for the clear-sky coefficients) or for which data is present and the calculation of cloudy radiances is requested through switching the flags `LCLD_RTCALC_SCREEN` or `LCLD_RTCALC_ASSIM` on in the module `SATS_MIX` for the respective sensor.

The transmittance coefficient file for different sensors can use any number of fixed pressure levels, and the number can be different for different sensors. However, interpolation to RTTOV pressure levels is performed inside RTTOV - see 3.3.2 - so the IFS does not need to know about that.

Prior to 50r1, RTTOV did not simulate the visible channels and infrared-only optical depth coefficient files were used for geostationary sensors. As the channel numbering in RTTOV coefficient files always begins at one, with this configuration, we used the channel numbering convention from the RTTOV coefficient files (e.g., in SEVIRI Infrared-only coefficient files, channel 2 was 6.3 micron and channel 3 was 7.4 micron).

New Visible/Infrared coefficient files including all channels replaced infrared-only optical depth coefficient files for geostationary sensors in 50r1. Accordingly, the channel order in the RTTOV coefficient files now matches the convention used for each instrument (e.g., in SEVIRI Visible/Infrared coefficient file channel 5 is 6.3 micron and channel 6 is 7.4 micron). Technical updates to geostationary satellite channel numbering have been accounted for in the geostationary radiances processing path to ensure the system uses from now onward the instrument channel numbering. This is equivalent with an offset of 3 for SEVIRI and an offset 6 for GOES and Himawari.

3.3.2 Clear-sky nadir radiances

(a) *Bias correction*

The bias correction is performed through variational bias correction, see part II. The VarBC class is called "rad", and the class-specific routines for the generic VarBC code are in the module `VARBC_RAD`. For most sounding radiances, the predictors used consist of four layer-thicknesses derived from the First Guess (as defined in `VARBC_PRED`), but some window channels do not include such airmass predictors to avoid aliasing of cloud information into the bias correction. The highest-peaking microwave

temperature channel (AMSU-A 14, ATMS 15, and MWS 16) is assimilated with a constrained bias correction, in order to better anchor the stratospheric temperature analysis (Han and Bormann, 2016). Without such an anchor, the variational bias corrections tend to drift to unrealistic values as a result of model biases. This is done at script level, through the namelist NAMVARBC_RAD.

(b) *Calling the radiative transfer model*

The radiative transfer model RTTOV is called from the general observation operator routine `HOP` via `OBSOP_RAD`, `RADTR_ML` and `RTTOV_EC`. More details on RTTOV can be found in Eyre (1991) and Saunders *et al.* (2018).

RTTOV performs optical depth calculations on a number of fixed pressure levels, as specified in the `rtcoef_` coefficient file mentioned earlier. RTTOV includes the option to provide the atmospheric profile input either on this set of fixed pressure levels or on a set of different and variable pressure levels. In the latter case, RTTOV will perform the required interpolation internally, using an interpolation that provides smoother gradients than a simple linear interpolation. The IFS uses this option, so that the radiative transfer computations are performed on the input user levels, which are the NFLEVG (currently 137) IFS full levels (plus pressure boundaries on NFLEVG + 1 half levels) rather than the fixed RTTOV pressure levels.

The vertical interpolation to the RTTOV pressure levels is now performed by the RTTOV internal interpolation (Hocking, 2014). This was previously done using the same PP routines involved in conventional observation processing, but this caused a problem of ‘missing levels’ in the TL and adjoint. Because of the limited number of RTTOV coefficient levels compared to IFS model levels, some IFS model levels were not used in the vertical interpolation and for these, the gradient of the model values with respect to the observations was zero, leading sometimes to vertical oscillations in the temperature increments. RTTOV internal interpolation is more sophisticated and ensures that all model levels are used in the vertical interpolation, preserving adjoint sensitivity at every level.

A check against the validity bounds of the RTTOV transmittance parametrization is in this case done within RTTOV (routine `SATRAD/RTTOV/MAIN/RTTOV_CHECKINPUT.F90`), for the optical depth calculations only (switch `APPLY_REGRESSION_LIMITS` in `SATRAD/MODULE/RTTOV_CONST.F90`). Various radiance preparations are again performed in `HRADP_ML`.

In either case, the routine `OBSOP_RAD` constructs a list of requested channel numbers for each report from the observation array, and only model radiances for exactly those channels are then requested from the RT-code. The routine `RADTR_ML` packets the profiles into chunks of work of an appropriate maximum size for the RT-code (currently set to 8 in `SATRAD/MODULE/mod_cparam.F90`). The RT packet size has been communicated to IFS in the call to `RTSETUP`. The output is radiances for the channels requested.

The TL and the adjoint are done in the common routines `HOP` and `OBSOP_RAD`, with separate lower-level routines. The \dot{T} and \dot{q} have to be recomputed before the actual tangent linear and adjoint computations can start. The pointers to the radiance data in the observation array are obtained just as done in the direct code. Consistency of the TL and adjoint operator can be tested by turning on the switch `LRTTOV_ADTEST` in prepIFS.

(c) *Skin-temperature ‘sink-variable’ at satellite FOVs*

In the case of 1C, or ‘raw’ radiance data, as used since May 1999 (McNally *et al.*, 1999) surface skin temperature is retrieved by 3D/4D-Var at each field of view, if the switch LTOVSCV is on (default is on), in namelist NAMVAR. This is done for all infrared and microwave satellite sensors and instruments. The handling of the skin temperature retrieval at the radiance FOVs is performed in the routine `HRADP/HRADPTL/HRADPAD`, called from `OBSOP_RAD`. The background skin temperature is provided by the model trajectory integration, and a background error of 1 K/5 K/7.5 K is assigned over sea/land/sea-ice, respectively (set in `SURAD`). The gradient with respect to the skin temperature obtained from RTTOV is temporarily stored in the TOVSCVX array and later transferred to its location in the distributed control vector. The next iteration of the minimisation provides updated skin temperature

increments (also stored in TOVSCVX) that are used by RTTOVTL in subsequent iterations. The outer-loop iterations result in a new linearisation state, stored in TOVSCVX5. All the skin-temperature-related information at FOV locations that needs to be passed between job-steps, reside in the ODB, in the skintemp array. Here, skintemp.1 is the background skin temperature, and subsequent values are the values at the end of each following minimisation. The approach has been adopted for CO₂ retrieval at AIRS FOVs (Engelen *et al.*, 2004).

3.3.3 All-sky nadir radiances

Observations from microwave imagers and sounders can be assimilated using an all-sky approach which unifies clear-sky, cloudy, and precipitation-affected radiances in one observation operator, using RTTOV for the radiative transfer. RTTOV is capable of modelling the effects of multiple scattering from hydrometeors. Three classes of microwave data are addressed by the all-sky route, with sometimes more than one class of data in the same instrument (e.g. AWS, MWS):

- (i) Microwave imagers—AMSR2, GMI, and some channels of SSMIS are actively assimilated; AMSU-A channels 1 and 2 are assimilated over ocean
- (ii) Microwave humidity sounders—MHS, MWHS-2 11-15, AWS 11-18, and some channels of SSMIS and GMI are actively assimilated over land, ocean and sea-ice
- (iii) Microwave temperature sounders—AMSU-A 5-14, MWHS-2 2-7, AWS 4-7 are assimilated over land, ocean and sea-ice with few exceptions

All-sky observations follow a path through the IFS code that is slightly different from that for clear-sky radiance observations, though the observation operator runs under **HOP** and **OBSOP_RAD** rather than under the model physics, as previously. However, the observation operator requires a number of diagnostic variables (e.g. precipitation flux and fraction) that come from the moist physics parametrizations and require a model timestep to have been run before they can be made available in observation space. These fields are generated in **CALLPAR**, stored in the **GFL** arrays and interpolated to observation space like any other model field required by an observation operator, using the **GOM** arrays. The main difference is that there is no interpolation: the observation operator gets the model profile at the grid point closest to the observation. An exception to this is temperature and humidity sounders such as AMSU-A and MHS, which use slant-path radiative transfer (Bormann, 2017) via 2D **GOM** arrays.

All-sky assimilation can be switched on in PrepIFS by means of switches for individual sensors, e.g. LAMSR2, LGMI, LSSMIS etc. For all-sky ATMS, which can run in parallel to clear-sky assimilation, there is a separate switch LATMS_ALLSKY. Comprehensive scientific documentation of the all-sky approach can be found in Bauer *et al.* (2010), Geer *et al.* (2010), Geer and Bauer (2011), Geer *et al.* (2014), Geer *et al.* (2022) and Duncan *et al.* (2022).

(a) Observation operator

Code for the all-sky observation operator is prefixed by 'mwave' and located either in **IFS/MWAVE** or **SATRAD/MWAVE**. The observation operator wrapper **IFS/MWAVE/mwave_wrapper** is called from **IFS/OP_OBS/OBSOP_RAD**. The same wrapper function is called whether in direct, TL or AD mode, and the required TL or AD functionality is driven by optional arguments. Lower down, there are separate observation operator routines for each mode: (**IFS/MWAVE/mwave_obsop**, **mwave_obsop_tl** or **mwave_obsop_ad**). In the screening trajectory, **IFS/MWAVE/mwave_screen** is called instead. Inputs to the observation operator are the profiles of model variables (passed in via a structure of **mwave_phys_type**) and any observation-related information (passed via a structure of **mwave_rad_type**). The observation-related information has been read from the ODB by **IFS/MWAVE/mwave_get**, **mwave_get_tl** or **mwave_get_ad**. Outputs from the observation operator (such as the simulated brightness temperatures, and the observation error) are returned in a similar way and are written to the ODB by **IFS/MWAVE/mwave_put** or **mwave_put_tl**. This approach is slightly different to that of other observation operators and is a result of the code's former location in the model physics.

The main function of the observation operator code in `IFS/MWAVE` and `SATRAD/MWAVE` is simply to provide the correct inputs and initialisations to run `RTTOV_SCATT`. However, there is code for quality control, to produce diagnostic output, and to determine the observation error, which is not constant, but a function of hydrometeor amount, as described in [Geer and Bauer \(2011\)](#). `VarQC` (see part II), `VarBC` (see part II), blocklisting and background quality control (all in [Chapter 4](#)) are done largely as for other radiance observations. However, some additional quality control and screening is sensor and channel specific. This includes screening of observations affected by Radio Frequency Interference (RFI), sun glint and Cold Air Outbreaks (see [Lonitz and Geer \(2015\)](#) and [Scanlon *et al.* \(2024\)](#) for further details).

A number of initialisation tasks are performed in `IFS/MWAVE/mwave.setup`, including reading the namelist `NAMMWAVE` for configuration flags. An array of `mwave.ids` structures is created, one for each satellite and sensor combination that will pass through the all-sky operators. In this private table are stored things like the instrument zenith angle, observation error specifications, and the ID numbers used in the rest of the IFS (e.g. sensor, satellite and bufr IDs).

As part of the all-sky observation operator, a set of ‘sink-variables’ are active which allow both the ocean skin temperature and the sea-ice concentration to be updated at each outer loop. These updated values are stored in `sink1.i` where `i` denotes the outer loop number with `sink1.1` storing the background values. Further information can be found in [Geer \(2024b\)](#), [Geer \(2024a\)](#) and [Scanlon *et al.* \(2024\)](#).

(b) External files

Observation error definitions for completely clear and completely cloudy skies are stored in files with names like `mw_err_< satellitseries >_< satellitenumber >_< instrumentname >`. These files define the type of observation error model (primarily, the choice of cloud proxy variable) to be used for each of a flexible number of different situations, currently land and sea. These files are read by `IFS/MWAVE/mwave.setup`. If the observation errors assigned have changed over time for a given sensor, these filenames have a timestamp appended to signify when the new errors became operational.

Some configuration options are specified in the `NAMMWAVE` namelist.

(c) Diagnostics

The consistency of the TL and adjoint operators can be tested by setting `ldmwave.test = .true.` in the `NAMMWAVE` namelist. This causes `IFS/MWAVE/mwave_obsop_test` to be called during the minimisation with the real TL input values for each observation. The results of the test are written to the IFS logfiles, prefixed by ‘MWAVE_OBSOP_TEST AD:’. The first number gives the TL / AD inaccuracy in multiples of machine precision. Typically this should be substantially less than 100 but it will go over 1000 for a few observations.

A number of diagnostics are stored in the ODB in the `allsky` or `allsky_body` table:

`DATUM.TBINFO`, `DATUM.TBFLAG`, `DATUM.TBERROR`, - These are bitfields which record processing and quality control decisions for the all-sky observations. These are an additional diagnostic on top of the usual status and event flags, and they only record decisions made internally in the all-sky observation operator. `DATUM.TBINFO` is purely informational, such as recording whether the observation was processed using the land path, open ocean (SST sink variable), cold ocean (sea ice sink variable) or other processing information. In contrast, if any bit in `DATUM.TBFLAG` is set, the datum (channel) is rejected and not included in data assimilation. `DATUM.TBERROR` is similar, but if any bit is set, no background departure is computed. These errors are where it is physically impossible to run the forward model (such as no valid surface emissivity estimate available, or `RTTOV` runtime errors) or where the observation itself is unavailable. Within this last bitfield, some bits represent an even higher class of errors where the whole observation must be rejected from data assimilation.

Binary arithmetic can be used to decipher these bitfields. For example, if `DATUM.TBFLAG = 33 = 25 + 20`, that means bits 5 and 0 have been set. Bit 5 indicates contamination by sea-ice. Bit 0 indicates coastal contamination. The full structure of the three `DATUM.TB` bitfields is described in `IFS/MODULE/yommwave`.

However, even if the observation or datum is considered OK by the all-sky observation operator (DATUM_TBFLAG and DATUM_TBERROR integer values are zero) it may subsequently be rejected by the other IFS screening processes (e.g. blocklisting, thinning, VarQC, background QC), so always check the 'status@hdr' and 'status@body' too.

DATUM_TBCLLOUD - this is a bitfield recording the status of diagnostic cloud and rain identification tests performed on observed and simulated brightness temperatures by IFS/OP_OBS/mwimager.cloud and further documented in the code and in Geer *et al.* (2008). The bitfield structure is documented in IFS/MODULE/yommwave. The lowest 2 bits give the results of the FG cloud test. With AND representing the bitwise boolean operator, (DATUM_TBCLLOUD AND 2) / 2 will give the result of the test for cloud in the FG, with 1 indicating a cloudy scene. (DATUM_TBCLLOUD AND 1) / 1 will give the result for the observation.

There are also a number of diagnostics relating to FG and analysis model state. These are valid at the time and location of the observation, giving information that is not otherwise archived. These values include the surface rain and snow rate, in $\text{kg m}^{-2} \text{s}^{-1}$, and the total columns of water vapour, cloud water, cloud ice, rain and snow, in kg m^{-2} .

3.3.4 Clear-sky limb radiances

Assimilation of clear-sky limb radiances has been implemented in the IFS for experimental purposes. The radiances are assimilated using the RTLIMB radiative transfer model which is an extension of RTTOV to the limb geometry. Details of the radiative transfer model and the assimilation of limb radiances can be found in Bormann *et al.* (2005); Bormann and Healy (2006); Bormann and Thépaut (2007); Bormann *et al.* (2007). Many aspects have been primarily developed for the assimilation of MIPAS limb radiances; the assimilation of radiances from other sensors is likely to require additional coding.

Limb radiances fall under obstype 10 "Limb observations", codetype 251. The general approach mirrors that used for clear-sky nadir radiances, i.e., the assimilation uses spatially interpolated vertical profile(s) of model variables. However, setup routines, the radiative transfer code, and the interface routines are different from the clear-sky nadir radiance assimilation.

The routine SULIMB sets up a limb group table (defined and stored in module YOMLIMB), for each satellite, sensor, and codetype, along the lines of the satellite group table set up in SURAD. Limb radiances are treated together with GPS radio occultation observations here. SULIMB calls the routine RTL_SETUP which includes the interface routine to the satrad library to read the RTLIMB coefficient files. Note that in contrast to the setup for nadir radiances, the code is set up to use fixed pressure levels, reference and limit profiles that are specific to the RT-coefficient files. This information is stored in the limb group table. RTL_SETUP also reads a channel selection file into the channel selection structure Y_LIMB_CHAN_SEL in the module YOMLIMB. Observation errors and constant, channel-specific biases are also read here from auxiliary files and stored in dedicated variables YOMLIMB.

Setting of observation errors and biases and screening of clear-sky limb radiances is performed from OBSOP_LIMB_RAD in the routines RTL_OBERROR and RTL_SCREEN. The latter applies the channel selection previously stored in Y_LIMB_CHAN_SEL in module YOMLIMB, and it performs cloud screening.

The actual assimilation happens under the routine HOP via OBSOP_LIMB_RAD. OBSOP_LIMB_RAD calls the routine RTL_HOP_1D which performs the following tasks: it interpolates the model profiles in the vertical to the fixed pressure levels (using the standard interpolation routines), does a simple extrapolation above the model top if required (based on a fixed mesospheric lapse rate for temperature, and holding humidity or ozone constant), and it checks the model profiles against the validity limits provided with the RTLIMB coefficient file. Finally, RTL_HOP_1D calls RTLIMB_HAT to enter the satrad library and perform the radiance computations.

The routine RTL_HOP_1D is used when local horizontal homogeneity is to be assumed for the radiative transfer computations. Alternatively, the radiative transfer computations can take the horizontal structure in the limb-viewing plane into account by providing a series of profiles covering the limb-viewing plane. In this case, profiles provided by 2D GOMs are used, and the routine RTL_HOP_2D is

called from `HOP` instead of `RTL_HOP_1D`. The two-dimensional facility is switched on by specifying `NOBSPROFS(10) > 1` for obstype 10 in the namelist `NAMNPROF`. Note that this means GPS radio occultation bending angles present in the assimilation will also take horizontal structure into account.

3.4 GPS RADIO OCCULTATION BENDING ANGLES

The subroutine `GPSRO.OP` is called from `OBSOP.GPSRO` and it simulates GPS radio occultation bending angles using the one-dimensional model outlined in Healy and Thépaut (2006). The subroutine evaluates a profile of bending angles, α as function of impact parameter, a , at each observation location by evaluating the integral

$$\alpha(a) = -2a \int_a^\infty \frac{\frac{d \ln n}{dx}}{(x^2 - a^2)^{1/2}} dx \quad (3.17)$$

where n is the refractive index and $x = nr$, the product of the refractive index and r , a radial coordinate value. The pressure, temperature, specific humidity and geopotential on the model levels (`ZPRESF5`, `ZTF5`, `ZQF5` and `ZGEOF5`, respectively) are the inputs to `GPSRO.OP`. The observation operator calculates the refractivity (defined as $N = 10^{-6}(n - 1)$) on the full model levels using the pressure, temperature and specific humidity profiles. It then converts the geopotential heights to geometric heights and then radius values. The bending angle integral is evaluated assuming that the refractivity, N , varies exponentially between the model levels.

The bending angle observation errors are set in `GPSRO_OBERROR`. Entries in the JO tables are set in `SULIMB`.

The routine `GPSRO.OP` is used when local horizontal homogeneity is to be assumed for the bending angle computations. Alternatively, the ray-tracing can take the horizontal structure in the limb-viewing plane into account by providing a series of profiles covering the limb-viewing plane. In this case, profiles provided by 2D-GOMs are used, and the routine `GPSRO_2DOP` is called from instead of `GPSRO.OP`. The 2d facility is switched on by specifying `NOBSPROFS(10) > 1` for obstype 10 in the namelist `NAMNPROF`. Note that this means any limb radiances present in the assimilation will also take horizontal structure into account.

3.5 GROUND-BASED RADAR PRECIPITATION COMPOSITES

This enables the assimilation of NEXRAD rain radar and rain gauges. The model-equivalent precipitation accumulations are computed inside the model physics and then stored directly to ODB. Here, the observation operator first accumulates the model surface precipitation fields at each time step over the same period (`NPRACCL`) as the observed precipitation composites. This accumulation is currently performed inside the model physics (`EC_PHYS`). Model precipitation amounts are then converted to $\log(\text{RR}[\text{mm/h}]+1)$ space, to be consistent with observations (in `GBRAD_OBSOP`). Scientific details can be found in Lopez (2011).

3.6 ATMOSPHERIC COMPOSITION

No documentation currently available.

3.7 CLOUD RADARS AND LIDARS

Observations from nadir-pointing cloud radar and lidar can be assimilated using the `ZmVar` observation operator (?). For radar, `ZmVar` has the capability to simulate both attenuated radar reflectivity and Doppler velocity. For lidar, `ZmVar` can simulate either attenuated total backscatter, attenuated molecular backscatter and cloud extinction.

The following instruments are currently supported:

- (i) EarthCARE CPR (94 GHz)—radar reflectivity (monitoring and assimilation) and Doppler velocity (currently monitoring only).

- (ii) EarthCARE ATLID (355 nm)—lidar extinction (currently monitoring only), total backscatter (currently monitoring only) and Rayleigh backscatter (currently monitoring only).
- (iii) CloudSat (94 GHz)—radar reflectivity (monitoring and assimilation).
- (iv) CALIPSO (532 nm)—lidar backscatter within cloud (monitoring and assimilation).

Radar and lidar observations follow a path through the IFS code that is very similar to the all-sky path. The observation operator is called from `HOP` and `OBSOP_RADLID`. As required for the simulation of all-sky radiances, precipitation fluxes from a call to the moist physics parametrizations are needed as input. These fields are generated in `CALLPAR`, stored in the `GFL` arrays and interpolated to observation space like any other model field required by an observation operator, using the `GOM` arrays. The observation operator uses the model profile at the grid point closest to the observation.

Radar and lidar assimilation can be switched on in PrepIFS using the switches `LCLRAD` and `LCLLID` respectively. Full details of the observation operators and observation handling are described in ?.

3.7.1 Observation operator

Code for the cloud radar and lidar observation operator is prefixed by 'clradlid' and located in `IFS/CLRADLID`. The observation operator wrapper `IFS/CLRADLID/clradlid_wrapper` is called from `IFS/OP_OBS/OBSOP_RADLID`. The same wrapper function is called whether in direct, TL or AD mode, and the required TL or AD functionality is driven by optional arguments. Lower down, there are separate observation operator routines for each mode (`IFS/CLRADLID/clradlid_obsop`, `IFS/CLRADLID/clradlid_obsop_tl` and `IFS/CLRADLID/clradlid_obsop_ad`). In the screening trajectory, `IFS/CLRADLID/clradlid_screen` is also called. Inputs to the observation operator are the profiles of model variables (passed in via a structure of `clradlid_phys_type`) and any observation-related information (passed via a structure of `rad_clref_type` or `lid_clbsec_type`). The observation-related information has been read from the ODB by `IFS/CLRADLID/clradlid_get`, `IFS/CLRADLID/clradlid_get_tl` or `IFS/CLRADLID/clradlid_get_ad`. Outputs from the observation operator (such as the simulated observations and the observation error) are returned in a similar way and are written to the ODB by `IFS/CLRADLID/clradlid_put` or `IFS/CLRADLID/clradlid_put_tl`.

A number of initialisation tasks are performed in `IFS/CLRADLID/clradlid_setup`, including reading the namelist `NAMCLRADLID` for configuration flags.

3.7.2 Observation error, screening and bias correction

The observation errors due to instrument error, forward model error and representativity error are all calculated separately using an 'error inventory' approach in `IFS/CLRADLID/clradlid_bc_err_screen`. Under the hypothesis of uncorrelated errors, the individual components are added in quadrature to define the total observation error. Due to the narrow field of view of the space-borne cloud radar and lidar observations, an appropriate definition of the representativity error is required to account for the mismatch of scales between the observations and the model, as described in [Fielding and Stiller \(2019\)](#).

The quality control for the observations is based on first-guess departures, thresholds for indicators of signal strength, estimated total attenuation and, for the radar, expected multiple scattering. The bias is based on a climatology of FG departures when the observations were passively monitored. To provide an implicit regime dependence for the correction scheme, indicators of temperature and height are used by the scheme.

Chapter 4

Screening

Table of contents

- 4.1 Introduction**
- 4.2 The structure of the observation screening**
 - 4.2.1 The incoming observations
 - 4.2.2 The screening run(s)
 - 4.2.3 General rationale of the observation screening
 - 4.2.4 3D-Var versus 4D-Var screening
- 4.3 Screening decisions made in observation operators**
 - 4.3.1 Satellite radiances
 - 4.3.2 Ground-based radar precipitation composites
- 4.4 Generic independent observation screening decisions**
 - 4.4.1 Preliminary check of observations
 - 4.4.2 Blocklisting
 - 4.4.3 Background (or rather trajectory) quality control
- 4.5 The dependent observation screening decisions**
 - 4.5.1 Update of the observations
 - 4.5.2 Global time–location arrays
 - 4.5.3 Vertical consistency of multilevel reports
 - 4.5.4 Removal of duplicated reports
 - 4.5.5 Redundancy check
 - 4.5.6 Thinning
 - 4.5.7 Compression of the ODB
- 4.6 Parallel aspects**
- Appendix A**
 - A.1 Bad reporting practice of SYNOP and TEMP reports
 - A.2 Revised background quality control for selected observations

4.1 INTRODUCTION

This chapter describes the observation screening in the ECMWF 3D/4D-Var data assimilation. A more general description can be found in [Järvinen and Undén \(1997\)](#). The purpose of the observation screening is to select a clean array of observations to be used in the data assimilation. This selection involves quality checks, removal of duplicated observations, thinning of their resolution etc.. The current selection algorithm has been operational since September 1996 and was to a large extent designed to reproduce the functionalities of the corresponding codes in the ECMWF OI analysis ([Lönnerberg and Shaw, 1985, 1987](#); [Lönnerberg, 1989](#)). Figure 4.1 illustrates the dataflow during the screening.

4.2 THE STRUCTURE OF THE OBSERVATION SCREENING

4.2.1 The incoming observations

The ‘extended’ ODB data base (the ECMA) contains all the observational information for the data window as required for 3D/4D-Var as well as all data that are going to be monitored. The next step

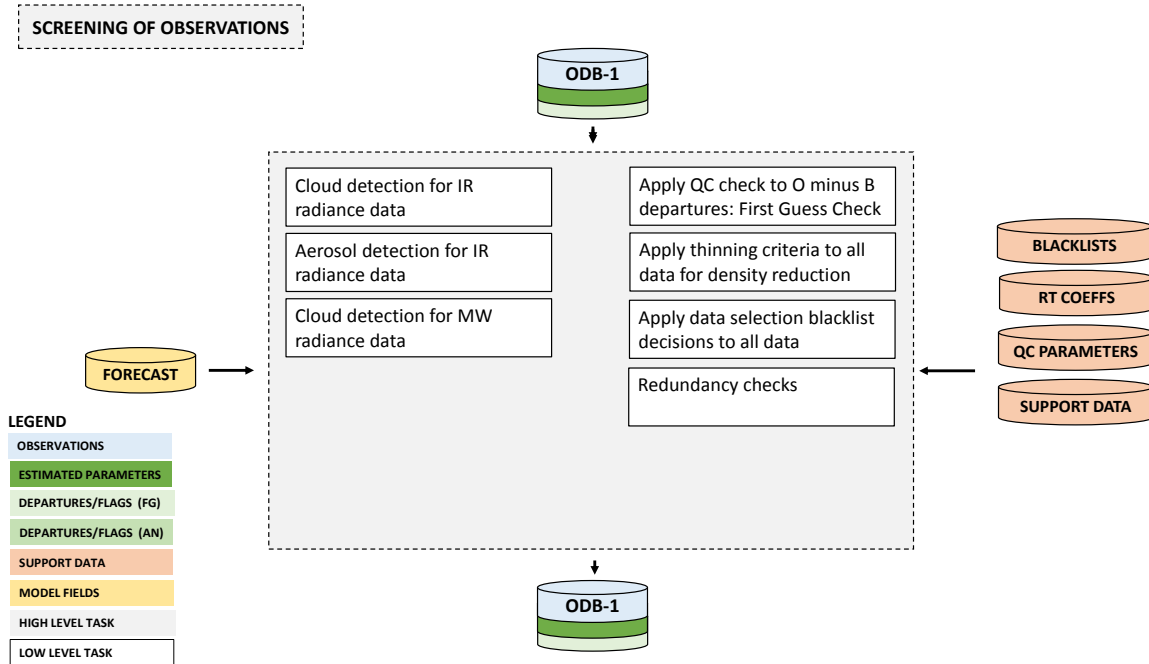


Figure 4.1 Simplified IFS flow diagram for screening. This is an illustrative, rather than comprehensive, description of the processing. See Fig. 1.1 for the wider observation processing context.

is that the observations are compared to the model as it is integrated for the length of the assimilation window. The observation minus model differences (the departures) are computed as described in Part II and stored in the ODB. These departures are an important input to the data selection procedures as many quality-control decisions depend on the magnitude of the departure. The collection of routines that perform data selection are jointly referred to as ‘the screening’. The purpose of the observation screening is to select the best quality observations, to detect duplicates, and reduce data redundancy through thinning.

4.2.2 The screening run(s)

The ECMWF 3D/4D-Var data assimilation system makes use of an incremental minimization scheme, as described in Part II. The sequence of jobs starts with the first (high resolution) trajectory run. During this run the model counterparts for all the observations are calculated through the non-linear observation operators, and the observation minus model difference (the departures) are calculated. As soon as these background departures are available for all observations, the screening can be performed. Prior to the screening the model fields are deallocated (dealmod) as most of the information necessary in the screening is stored in the observation data base (ODB). For the observation screening, the background errors (available as grid data in the ‘errgrib’ file, see Part II) are interpolated to the observation locations for the observed variables (**INIFGER**, **SUFGER** and **GEFGER** - see Chapter 2, section 2.6).

Technically, the final result of the observation screening is a pair of ODBs. The original ‘extended’ observation data base now contains observations complemented by the background departures, together with quality control information for most of the observations. This ECMA ODB remains on disc for later use in feedback creation. The compressed ODB, the CCMA, is a subset of the original observations, and is passed for the subsequent minimization job. The CCMA contains only those observations that are to be used in the minimization.

For several cycles in the early delivery, and from cycle 50r1 also in the long window data assimilation, aspects of continuous data assimilation have been included that change the observation preparation and screening process. The change is whether the screening runs once (on the background trajectory) or runs every trajectory except the last, allowing for newly arriving observations to be incorporated into the later minimisations that were not available when the background trajectory ran. So, in cycle

50r1, in any operational-like configuration, the full preprocessing (BUFR, BUFR2ODB) and the screening (LSCREEN=.true.) is re-run for every trajectory. As a result, each trajectory creates a new ECMA and then the CCMA containing the subset of screened observations is used by only one minimisation.

4.2.3 General rationale of the observation screening

The general logic in the 3D/4D-Var observation screening algorithm is to make the *independent* decisions first, i.e. the ones that do not depend on any other observations or decisions (**DECIS**). One example is the background quality control for one observed variable. These can be carried out in any order without affecting the result of any other independent decision. The rest of the decisions are considered as mutually *dependent* on other observations or decisions, and they are taken next, following a certain logical order. For instance, the horizontal thinning of radiance reports is only performed for the subset of reports that passed the background quality control. Finally, the CCMA database is created for the minimization in such a way that it only contains the data that will be used.

4.2.4 3D-Var versus 4D-Var screening

In the original 3D-Var assimilation system the screening rules were applied once, for the complete set of observations spanning a six-hour period. In the early implementation of the 4D-Var assimilation system, the same data selection approach called ‘3D-screening’ was applied over the 6-hour long 4D-Var time window, which resulted in essentially the same screening decisions as in 3D-Var.

In summer 1997, a new screening procedure called 4D-screening was implemented that took into account the temporal distribution of the observations. The time window is divided into time-slots of typically half-hour length (15 minutes for the first and the last time slots). The 3D-screening algorithm was then applied separately to observations within each time-slot. This allowed more data to be used by 4D-Var, for instance, all messages from an hourly reporting station can now be used, whereas only one (closest to central time) would have been allowed by the redundancy check in the 3D-screening. The 4D-screening behaviour is activated by switch LSCRE4D; it is meant to be used in conjunction with time correlation of observation errors where appropriate, as explained in [Järvinen *et al.* \(1999\)](#) and in Part II. From Cy50r1, the thinning time slots are independent from the 4D-Var ones, with the default value is specified with THINSEC (30 minutes currently). Custom time slots for observation groups as used in routine **SCREEN** could be set by passing the corresponding variable to **YOMSCC** namelist.

4.3 SCREENING DECISIONS MADE IN OBSERVATION OPERATORS

Many screening decisions are easiest to make in the observation operator itself, and are performed under the condition LSCREEN=.TRUE.

4.3.1 Satellite radiances

(a) *Cloud and rain rejection for clear-sky observations*

For data assimilated through the clear-sky route in the IFS, observations contaminated by significant cloud or rain signals must be removed before being supplied to the 4D-Var minimization in the clear-sky assimilation scheme. Microwave radiances that are assimilated in all-sky conditions are not subject to cloud or rain rejection.

For clear-sky microwave radiances (e.g. ATMS) cloud/rain screening is performed in the routine **MW_CLEARSKY_SCREEN**. This requires observation departures and is presently called from **DEPARTURE_JO**. In addition to departures in window channels, it also uses retrievals of liquid water path or scatter indices derived from the observations.

For infrared radiances the test for clouds is done in routine **CLOUD_DETECT** for AIRS/IASI/CrIS and routine **HIRS_CLD** for HIRS. The former is based on the algorithm described in [McNally and Watts \(2003\)](#). The latter is described in [Krzeminski *et al.* \(2009b\)](#). In both cases the aim is to identify which infrared channels can be used in a particular scene and which must be rejected.

For both the microwave and infrared data, if an observation is rejected due to cloud contamination, the `contam_cld_flag` is set in the `datum_event1` field in the ODB. These observations will not influence any aspects of the analysis including bias parameter evolution in the VARBC.

(b) *All-sky radiances*

Most data selection decisions for all-sky radiances are made within the observation operator code under `MWAVE_SCREEN` and recorded in the ODB column `DATUM_TBFLAG` (see [Chapter 3](#)).

4.3.2 Ground-based radar precipitation composites

Some additional screening is performed inside the observation operator for the NCEP Stage IV surface precipitation observations (routine `GBRAD_SCREEN`). Both sea points and points likely to be affected by the occurrence of atmospheric ducting are flagged as rejected. Ducting which leads to the undesirable deflection of radar beams towards the ground is diagnosed from vertical gradients of atmospheric refractivity, as computed from model temperature and moisture fields in routine `GBRAD_REFRACTION` (Lopez, 2011). In addition, points with background departures exceeding 1.5 in $\log(\text{RR}[\text{mm/h}]+1)$ space are also discarded from the assimilation.

4.4 GENERIC INDEPENDENT OBSERVATION SCREENING DECISIONS

4.4.1 Preliminary check of observations

The observation screening begins with a preliminary check of the completeness of the reports (`PRECH`). None of the following values should be missing from a report: observed value, background departure, observation error and vertical coordinate of observation. Also a check for a missing station altitude is performed for SYNOP, TEMP and PILOT reports. The reporting practice for SYNOP and TEMP mass observations (surface pressure and geopotential height) is checked (`REPRE`), as explained in [Appendix A](#). At this stage also, the observation error for SYNOP geopotential observations is inflated if the reported level is far from the true station level (`ADDOER`). The inflation is defined as a proportion of the difference between the reported level and the true station altitude by adding 2% of the height difference to the observation error.

4.4.2 Blocklisting

Next, the observations are scanned through for blocklisting (still 'blacklisting' in the code, subroutine `BLACK`). At the set-up stage the blocklist interface is initialized (`BLINIT`) to the external blocklist library. The blocklist files consist formally of two parts. Firstly, the selection of variables for assimilation is specified in the 'data selection' part of the blocklist file. This controls which observation types, variables, vertical ranges etc. will be selected for the assimilation. Some more complicated decisions are also performed through the data selection file; for instance, an orographic rejection limit is applied in the case of the observation being too deep inside the model orography. This part of the blocklist also provides a handy tool for experimentation with the observing system, as well as with the assimilation system itself. Secondly, a 'monthly monitoring' blocklist file is provided for discarding the stations that have recently been reporting in an excessively noisy or biased manner compared with the ECMWF background field.

Most data selection criteria are coded in so called *blocklist files*, written in a convenient, readable blocklist language (see the Blocklist Documentation, [Järvinen et al., 1996](#)). The blocklist mechanism is very flexible and allows nearly complete control of which data to use/not use in the assimilation. The 'monthly blocklist' is the part of the blocklist that is based on operational data monitoring results, and it is maintained by the Meteorological Operations Section. The blocklist is consulted in the screening job. The interface is set up in `BLINIT`, in such a way that a number of named items from the header ([Table 4.1](#)) and body ([Table 4.2](#)) parts of the observation report can be passed to the blocklist software. Depending on the blocklisting criteria flags are communicated to the routine `BLACK`, and those are written to the ECMA ODB data base. Blocklist-rejected data are subsequently excluded from the CCMA ODB and will not be present in the minimisation job steps. Data selection rules should be coded in the blocklist

Table 4.1 Header variables in the ifs/blacklist interface. Exact contents may change - see *black.F90*.

Index	Name	Description
1	OBSTYP	observation type
2	STATID	station identifier
3	CODTYP	code type
4	INSTRM	instrument type
5	DATE	date
6	TIME	time
7	LAT	latitude
8	LON	longitude
9	STALT	station altitude
10	LINE_SAT	line number (atovs)
11	RETR_TYP	retrieval type
12	QI.1	quality indicator 1
13	QI.2	quality indicator 2
14	QI.3	quality indicator 3
15	MODORO	model orography
16	LSMASK	land-sea mask (integer)
17	RLSMASK	land-sea mask (real)
18	MODPS	model surface pressure
19	MODTS	model surface temperature
20	MODT2M	model 2-metre temperature
21	MODTOP	model top level pressure
22	SENSOR	satellite sensor indicator
23	FOV	field of view index
24	SATZA	satellite zenith angle
25	NANDAT	analysis date
26	NANTIM	analysis time
27	SOE	solar elevation
28	QR	quality of retrieval
29	CLC	cloud cover
30	CP	cloud top pressure
31	PT	product type
32	SONDE_TYPE	sonde type
33	SPECIFIC	amsua specific
34	SEA_ICE	model sea ice fraction

files whenever possible rather than in the IFS code itself. The operational blacklist history is kept in an archive.

(a) *Scatterometer blocklisting decisions*

In order to screen on sea-ice contamination, scatterometer data are removed (within the blacklist mechanism) whenever the model sea-ice fraction exceeds 1% or the model sea-surface-temperature analysis is below 273.15 K. Land is removed by imposing that the model land-sea mask should not exceed 10%.

(b) *Satellite radiances*

Like for any other observations, decisions are made to use or not use a particular radiance observation in the blacklist. These fall into two distinct types: The first is the usual a priori type decision which takes no account of the actual value of the observation. Examples for radiances include the exclusion of data measured by new instruments which we do not yet wish to use, data measured by bad/failed instruments, data measured at extreme scan positions, exclusion of data measured over land or high orography and the exclusion of data at certain times of year when solar intrusions may cause problems

Table 4.2 Body variables for the ifs/blacklist interface. Exact contents may change - see [black.F90](#)

Index	Name	Description
1	VARIAB	variable name
2	VERT_CO	type of vertical coordinate
3	PRESS	pressure, height or channel number
4	PRESS_RL	reference level pressure
5	PPCODE	SYNOP pressure code
6	OBS.VALUE	observed value
7	FG_DEPARTURE	first guess departure
8	OBS-ERROR	observation error
9	FG_ERROR	first-guess error
10	WINCHAN_DEP	window channel departure
11	OBS.T	observed temperature

(there are others). The second type of test is particular to radiances and is a run-time decision based on the observed values (or more correctly the radiance departure from the background).

Depending on the magnitude of the radiance departure in key window channels, individual or combinations of microwave and infrared channels may be rejected. In some respects this may be considered an additional first-guess check that takes place in the blocklists. It can equally well be considered as an additional cloud/rain detection check that takes place in the blocklist as it exclusively involves window channels. No attempt is made here to document the particular test and threshold which are applied to each channel on every instrument and the user is referred to the data selection files for details. For both types of test applied in the blocklists environment, if it is failed there are two options for what then results. The setting of a FAIL(CONSTANT) flag means that the observations will be rejected and take no further part in the analysis. The setting of a FAIL(EXPERIMENTAL) flag means that the observation will enter the main analysis in such a way that it cannot force increments of e.g. temperature or humidity, but it can influence the calculation and evolution of bias correction coefficients inside VARBC. An example of when the latter is used would be for a new satellite for which we do not wish to actively assimilate the data, but wish to establish an accurate bias correction. Another application of the FAIL(EXPERIMENTAL) facility is its use for window channels used in the quality control of other data.

For some window channel microwave radiances over land, another setting of the blocklisting decision can be FAIL(USE_EMISKF_ONLY). This means that the observation will not be used for the atmospheric analysis in a similar way as FAIL(CONSTANT) rejects the observation, but the observation is used to influence the emissivity Kalman Filter atlas, described in section 3.3.2.

4.4.3 Background (or rather trajectory) quality control

The background quality control (FIRST) is performed for all the variables that are intended to be used in the assimilation. The procedure is as follows. The variance of the background departure $y - H(\chi_b)$ can be estimated as a sum of observation and background-error variances $\sigma_o^2 + \sigma_b^2$, assuming that the observation and the background errors are uncorrelated. After normalizing with σ_b , the estimate of variance for the normalized departure is given by $1 + \sigma_o^2/\sigma_b^2$. In the background quality control, the square of the normalized background departure is considered as suspect when it exceeds its expected variance more than by a predefined multiple ([FGCHK](#), [SUFGLIM](#)). For the wind observations, the background quality control is performed simultaneously for both wind components ([FGWND](#)). In practice, there is an associated background quality-control flag with four possible values, namely 0 for a correct, 1 for a probably correct, 2 for a probably incorrect and 3 for an incorrect observation, respectively ([SUSCRE0](#)). [Table 4.3](#) gives the predefined limits for the background quality control in terms of multiples of the expected variance of the normalized background departure. These values are set in DEFRUN and can be changed in namelist [NAMJO](#). For SATOB winds the background error limits are modified as explained in [Appendix A](#).

Table 4.3 The predefined limits for the background quality control, given in terms of multiples of the expected variance of the normalized background departure.

Variable	Flag 1	Flag 2	Flag 3
u, v	9.00	16.00	25.00
z, ps	12.25	25.00	36.00
dz	x	x	x
T	9.00	16.00	25.00
rh, q	9.00	16.00	25.00

Flag values are denoted by 1 for a probably correct, 2 for a probably incorrect and 3 for an incorrect observation. The variables are denoted by u and v for wind components, z for geopotential height, ps for surface pressure, dz for thickness, T for temperature, rh for relative humidity and q for specific humidity, respectively.

There is also a background quality control for the observed wind direction (**FGWND**). The predefined error limits of 60°, 90° and 120° apply for flag values 1, 2 and 3, respectively. The background quality control for the wind direction is applied only above 700 hPa for upper-air observations for wind speeds larger than 15 ms⁻¹. If the wind-direction background quality-control flag has been set to a value that is greater than or equal to 2, the background quality-control flag for the wind observations is increased by 1.

There is no first-guess check for scatterometer data. It is demanded, though, that neither scatterometer nor model wind speed should exceed 35 ms⁻¹, since that marks the range of validity for scatterometer wind inversion.

As with all the screening, the background check runs under `LSCREEN=true`. So under continuous data assimilation, i.e. rescreening, which is standard from 50r1, this means that it is no longer just a background check but a check based on the latest updated trajectory.

4.5 THE DEPENDENT OBSERVATION SCREENING DECISIONS

4.5.1 Update of the observations

Just before performing the dependent screening decisions, the flag information gathered so far is converted into a status of the reports, namely: **active**, **passive**, **rejected** or **blocklisted**, and also into a status of the data in the reports (**FLGTST**). The reports with a RDB report flag value 2 (probably incorrect) or higher for latitude, longitude, date and time are rejected. For the observed data there are RDB datum flags for the variable and for the pressure, i.e. the pressure level of the observation. The rejection limits for these are as follows: all data are rejected for the maximum RDB datum flag value 3 (incorrect), non-standard-level data are rejected for the maximum RDB datum flag value 2, and for the pressure RDB datum flag the rejection limit is 1 (probably correct). The background quality control rejection limits are flag value 3 for all the data, and flag value 2 for the non-standard-level data.

4.5.2 Global time–location arrays

Some of the dependent decisions require a global view to the data which is not available as the memory is distributed. Therefore *ad hoc* global time–location arrays are formed and broadcast in order to provide this view (`GLOBA`, `DISTR`).

4.5.3 Vertical consistency of multilevel reports

The first dependent decisions are the vertical-consistency check of multilevel reports (**VERCO**), and the removal of duplicated levels from the reports. The vertical-consistency check of multilevel reports is

applied in such a way that if four consecutive layers are found to be of suspicious quality, even having a flag value one, then these layers are rejected, and also all the layers above these four are rejected in the case of geopotential observations. These decisions clearly require the quality-control information, and they are therefore 'dependent' on the preceding decisions.

4.5.4 Removal of duplicated reports

The duplicated reports will be removed next. That is performed (**MISCE**, **DUPLI**, **REDSL**) by searching for pairs of exactly collocated reports of the same observation types, and then checking the content of these reports. It may, for instance, happen that a report is formally duplicated by having a slightly different station identifier but with the observed variables inside these reports being exactly the same, or partially duplicated. The pair-wise checking of duplicates results in a rejection of some or all of the content of one of the reports.

4.5.5 Redundancy check

The redundancy check of the reports, together with the level selection of multi-level reports, is performed next for the active reports that are collocated and that originate from the same station (**REDUN**). In 3D-screening, this check applies to the whole observation time window. In 4D-screening (**LSCRE4D** = **.TRUE.**), this check applies separately in each timeslot.

For **LAND SYNOP** reports, the report closest to the analysis time with most active data is retained, whereas the other reports from that station are considered as redundant and are therefore rejected from the assimilation (**REDRP**, **REDMO**). For **SHIP SYNOP** and **BUOY** observations the redundancy check is done in a slightly modified fashion (**REDGL**). These observations are considered as potentially redundant if the moving platforms are within a circle with a radius of 1° latitude. Also in this case only the report closest to the analysis time with most active data is retained. All the data from the multilevel **TEMP** and **PILOT** reports from same station are considered at the same time in the redundancy check (**REDOR**, **SELEC**). The principle is to retain the best quality data in the vicinity of standard levels and closest to the analysis time. One such datum will, however, only be retained in one of the reports. A wind observation, for instance, from a sounding station may therefore be retained either in a **TEMP** or in a **PILOT** report, depending on which one happens to be of a better quality. A **SYNOP mass** observation, if made at the same time and at the same station as the **TEMP** report, is redundant if there are any **TEMP** geopotential height observations that are no more than 50 hPa above the **SYNOP mass** observation (**REDSM**).

4.5.6 Thinning

Finally, a horizontal thinning is performed for the aircraft data, some radiances (**ATOVS**, **AIRS**, **IASI**, etc.), **GEOS**, **SATOB**, **ERS** and **SCAT** reports. The horizontal thinning of reports means that a predefined minimum horizontal distance between the nearby reports from the same platform is enforced.

Aircraft reports were originally thinned checking one aircraft at a time (**MOVPL**, **THIAIR**), with a tolerance of about 60 km between used reports ([Cardinali et al., 2003](#)). Vertical thinning criteria allowed the assimilation of profiles from ascending and descending flight paths. In 2022 it was found that this "track checking" was allowing excessive numbers of Mode-S aircraft reports to be used over central Europe. Late in 2023 there was a change to use "box thinning" for aircraft data (**NEW_THINN_AIRCRAFT**), similar to that used for satellite data below. the details are in [Ingleby \(2025\)](#).

At 49r2 "box thinning" was introduced for surface data - partly to cope with increasing numbers and complexity of reports especially from land stations. The top level routine (**ONETHIN_SURF**) sets up both the box numbers (using **THIBOX**) and a preference order based on observation type (**BUFR SYNOP** first, **SHIP** and **METAR** last), there is also a preference for station heights close to the model orography and for active surface pressure. A single sort/pass is performed with a box width of 0.15 degrees latitude, about 16.7 km. Following tests it was decided to thin both land and marine reports together.

In May 2025, together with the introduction of global Mode-S there was a further change to aircraft thinning with a new top level routine ([ONETHIN_AIRCRAFT](#)). This simplifies the aircraft thinning and allows preference by report type (AMDAR first, Mode-S last) to work as intended. Also based on diagnostic results from [Ingleby \(2025\)](#) the box size used for aircraft is made a function of pressure: 44 km at low levels rising to 95 km at 200 hPa.

Thinning of some radiances, GRAD, SATOB, ERS and ASCAT SCAT reports are each done in two stages controlled by THINN. For radiances ([THINNER](#)), a minimum distance of about 70 km is enforced and, thereafter, a repeat scan is performed to achieve the final separation of roughly 250 km or 120 km between reports from one platform. This is controlled through settings in [DEFRUN](#), that can also be modified through namelist ([NAMSCC](#)). The thinning algorithm is the same as used for AIREPs except that for radiances a different preference order is applied: a sea sounding is preferred over a land one, a clear sounding is preferred over a cloudy one and, finally, the closest observation time to the analysis time is preferred. For geostationary water vapour radiances, a similar thinning in two stages is applied with currently about 36 km minimum distance and about 72 km final separation ([THINNER](#)). During the thinning, preference is given to data having the largest fraction of clear sky in the clear-sky radiance average, high infrared brightness temperature (for GOES data) and, finally, a small standard deviation of brightness temperatures within the CSR mean. A similar thinning technique is applied to SATOB high-density data ([THINNER](#)). Note that prior to assimilation a coarser pre-thinning may take place already during observation pre-processing in order to reduce otherwise excessive data volumes.

The screening of SATOB data has been extended for atmospheric motion wind observations, including individual quality estimate. The quality information from the quality control performed by the producer at extraction time is appended to each wind observation. This Quality Indicator (QI) is introduced as an additional criterion in the thinning step; priority is given to the observation with the highest QI value.

For ASCAT scatterometer data, the above described thinning algorithm is only applied along track. In across-track direction, backscatter data from these platforms are provided into wind-vector cells (WVC) with a spatial resolution of 25 km. In this direction, data is thinned by selecting predefined wind-vector cells (subroutine [SCAQC](#)). For ASCAT, from 42 cells (two swaths of 21 cells each) the odd cells are used in the left swath and the even ones in the right swath. After this across-track thinning, the generic thinning algorithm is applied to the remaining cells in along-track direction. Overall, one observation every roughly 50 km is used in the across-track and along track directions. QuikSCAT data (also provided on a 25 km grid) are not thinned. Instead, a 50 km wind product is determined from backscatter data from four underlying 25 km cells, each given a reduced weight of one fourth.

ASCAT scatterometer winds are, besides thinning, subject to a high-wind rejection test with an upper-wind speed limit set to 35 ms^{-1} to both the scatterometer and background winds ([FGWND](#)).

No all-sky microwave observations are included in this thinning, since in most cases the observations being assimilated are superobservations, and any thinning that is done is done in the context of the equal-area superobbing grid (currently reduced Gaussian) and operates by keeping one in every N observation per row and only keeping every M rows. See the all-sky code for further info; this runs under MWAVE.SCREEN in the observation operator and is done prior to all the screening processes described in the current chapter.

4.5.7 Compression of the ODB

After the observation screening roughly a fraction of 1/10 of all the observed data are active and so the compressed observation ODB (the CCMA) for the minimization run only contains those data. The large compression rate is mainly driven by the number of radiance data, since after the screening there are only 10–20% of the radiance reports left, whereas for the conventional observations the figure is around 40%. As a part of the compression, the observations are re-sorted amongst the processors for the minimization job in order to achieve a more optimal load balancing of the parallel computer.

4.6 PARALLEL ASPECTS

As mentioned earlier, in the observation screening there are the two basic types of decision to be made. Independent decisions, on one hand, are those where no information concerning any other observation or decision is needed. In a parallel-computing environment these decisions can be happily made by different processors fully in parallel. For dependent decisions, on the other hand, a global view of the observations is needed which implies that some communication between the processors is required. The observation array is, however, far too large to be copied for each individual processor. Therefore, the implementation of observation screening at the ECMWF is such that only the minimum necessary information concerning the reports is communicated globally.

The global view of the observations is provided in the form of a global 'time–location' array for selected observation types. That array contains compact information concerning the reports that are still active at this stage. For instance, the observation time, location and station identifier as well as the owner processor of that report are included. The time–location array is composed at each processor locally and then collected for merging and redistribution to each processor. After the redistribution, the array is sorted locally within the processors according to the unique sequence number. Thus, every processor has exactly the same information to start with, and the dependent decisions can be performed in a reproducible manner independently of the computer configuration.

The time–location array is just large enough for all the dependent decisions, except for the redundancy checking of the multilevel TEMP and PILOT reports. This is a special case, in the sense that the information concerning each and every observed variable from each level is needed. Hence, the whole multilevel report has to be communicated. The alternative to this would be to force the observation clusters of the multilevel reports always into one processor without splitting them. In that case the codes responsible for the creation of the observation arrays for assimilation would need to ensure the geographical integrity of the observation arrays distributed amongst the processors. This is, however, not possible in all the cases, and the observation screening has to be able to cope with this. Currently, it is coded in such a way that only a limited number of multilevel TEMP and PILOT reports, based on the time–location array, are communicated between the appropriate processors as copies of these common stations.

APPENDIX A

A.1 Bad reporting practice of SYNOP and TEMP reports

The way the synoptic surface stations report mass observations (pressure or geopotential height) is considered as bad if:

- station altitude is above 800 m and station reports mean sea level pressure
- station altitude is above 800 m and station reports 1000 hPa level
- station altitude is above 1700 m and station reports 900 hPa level
- station altitude is below 300 m and station reports 900 hPa level
- station altitude is above 2300 m and station reports 850 hPa level
- station altitude is below 800 m and station reports 850 hPa level
- station altitude is above 3700 m and station reports 700 hPa level
- station altitude is below 2300 m and station reports 700 hPa level
- station altitude is below 3700 m and station reports 500 hPa level

The reporting practice is also considered as bad if the station reports 500 gpm, 1000 gpm, 2000 gpm, 3000 gpm or 4000 gpm level pressure, respectively, and station altitude is more than 800 m different from the reported level.

For TEMP geopotentials the reporting practice is considered as bad if:

- station altitude is above 800 m and station reports 1000 hPa level
- station altitude is above 2300 m and station reports 850 hPa level
- station altitude is above 3700 m and station reports 700 hPa level

A.2 Revised background quality control for selected observations

The background quality-control rejection limits are applied more strictly for some observation types than stated in [Table 4.3](#). The special cases are the following ones.

- AIREP wind observations with zero wind speed are rejected if the background wind exceeds 5 m s^{-1} .
- For AIREP and BUOY wind observations the rejection limit is multiplied by 0.5, and for PILOT wind by 0.8.
- For SATOB wind observations the rejection limit is multiplied by 0.1, except below 700 hPa level where it is multiplied by 0.2.
- No background quality control is applied for SCAT winds.
- For BUOY surface pressure observations the rejection limit is multiplied by 0.9.
- For AIREP temperature observations the rejection limit is multiplied by 1.6.

Chapter 5

Deprecated areas

Table of contents

[5.1 Introduction](#)

[5.2 Observation pre-processing](#)

5.1 INTRODUCTION

This section is for deprecated parts of the observation code, things which have increasingly little relevance to understanding the modern system. Unfortunately, the deprecated code still performs a few important tasks, otherwise it would have already been removed. This documentation is not kept up to date but it is still useful for understanding those few remaining tasks, especially for those people whose job it will be to replace them completely. There are at least two observation CODETYPEs that are not used operationally, but may be used by reanalysis: ROCOB (rocket observations) and PAOB (originally Pseudo-Australian Observations, used in the earlier decades of ERA5 for tropical cyclone position/pressure estimates). Unfortunately US wind-profiler data stopped in about 2014 (data from a few stations were received until 2017). Data formerly reported as ARARS or ACARSMR are now received as WIGOSAMDAR.

5.2 OBSERVATION PRE-PROCESSING

For several years most **MKCMARPL** functionality was replaced by COPE filters for conventional data. Because the COPE filters were less computationally efficient this change was reverted at 48r1. The **MKCMARPL** and associated code is now even older and remains an ideal candidate for a thorough rewrite.

A number of non-conventional observation types still rely on **MKCMARPL** in important ways, or just in annoying ways that make it hard to simply delete what is apparently useless code. For example:

- AMVs require variable conversion from windspeed and direction to horizontal wind components *u* and *v*. More pointlessly, they also need to set up an (unused) initial observation error using the deprecated **OBSERR** routines. This error needs to be present in the ODB to prevent rejection during screening, even though it is not used.
- For some satellite radiances, channel numbering is reassigned. There may or may not be anything else useful going on under **LEVEL1CGEOS.OB**. For example, there are checks on longitudes, latitudes and times, but these are likely duplicates of checks elsewhere in the pre-processing.
- Ozone observations rely on **MKCMARPL**. Among other things they set up observation errors here.

Chapter 6

Tables, codes and flags

Table of contents

- 6.1 Introduction
- 6.2 Obstype and codetype
- 6.3 Variable codes (varno)
- 6.4 Additional observation operator codes: NVAR and CVAR_NAME
- 6.5 Observation characteristics: instrument specification and retrieval type
- 6.6 Vertical coordinate: pressure, satellite ID and level ID codes
- 6.7 ODB report status: events, flags and codes
- 6.8 Datum status: events, RDB and analysis flags

6.1 INTRODUCTION

Older versions of the IFS documentation have contained copious tables attempting to document all the codes, flags and groupings for observations in the IFS. If this was ever possible, in recent years the variety of observational data processing has mushroomed, these tables have in some places become hopelessly out of date and their true counterparts are too large to comfortably fit in such documentation. The truest record is the IFS source code, the ODB data archived to MARS, and the logfiles of the IFS. The ODB governance database at <http://apps.ecmwf.int/odbgov> is reasonably up-to-date but does not cover everything. Some of the old tables from the IFS documentation are provided here, although no longer kept up-to-date, as they are useful for illustration purposes, and as a guide to what these codes are for. See also Sec. 1.4.

6.2 OBSTYPE AND CODETYPE

In the JO table results are presented by OBSTYPE and within that by CODETYPE. To confuse matters there are also SUBTYPE and REPORTTYPE, providing similar granularity to CODETYPE. (The reasons for these different sets of labels are historical. SUBTYPE is/was more closely aligned with the WMO code formats.) Some CODETYPES are listed in Table 6.1. See the ODB governance tables <http://apps.ecmwf.int/odbgov> for a fuller and more up-to-date list. The coexistence of different codes used for BUFR and ODB observation types and the subtype and code type requires a mapping from one to another. The ODB governance tables also describe that mapping. There are some minor CODETYPES, such as 'Abbreviated SHIP' and SHRED, that are almost obsolete. CODAR is obsolete, COLBA (constant level balloon) is rarely used. For Land SYNOP and radiosondes there are both Traditional Alphanumeric Code (TAC) and binary BUFR formats in use. Dropsonde reports are intermittent.

Table 6.1 Some ODB observation types and code types. See the ODB governance tables <http://apps.ecmwf.int/odbgov> for a fuller and more up-to-date list

OBSTYPE		CODETYPE	
Code	Name	Code	Name
1	SYNOP	11	Manual Land SYNOP (TAC)
		14	Automatic Land SYNOP (TAC)
		21	Manual SHIP (TAC)
		24	Automatic SHIP (TAC)
		170	BUFR Land SYNOP
		182	BUFR SHIP
		140	METAR
		110	GroundGPS
2	AIREP (Aircraft)	AIREP (TAC)	
		144	AMDAR (TAC/early BUFR)
		145	TAMDAR
		145	WIGOS AMDAR (BUFR)
		241	ADS-C (BUFR)
3	SATOB (AMV)	88	SATOB
		89	High Resolution VIS wind
		90	AMV
		188	SST
4	BUOY (aka DRIBU)	188	Moored Buoy (BUFR)
		189	Drifting Buoy (BUFR)
5	TEMP (Radiosonde)	35	Land TEMP (TAC)
		36	SHIP TEMP (TAC - rare now)
		109	BUFR Land TEMP
		111	BUFR Ship TEMP
		135	DROP TEMP (TAC)
		230	BUFR DROP TEMP
231	BUFR TEMP DESCENT		

Table 6.1 *Continued.*

Observation Type		Code Type	
Code	Name	Code	Name
6	PILOT	32	Land PILOT (TAC)
		112	BUFR Land PILOT
		131	Japanese Wind Profiler
		134	European Wind Profiler
7	SATEM	86	GTS SATEM
		184	High Resolution Simulated TOVS
		185	High Resolution Simulated DWL SATEM
		186	High Resolution SATEM
		200	GTS BUFR SATEM 250km
		201	GTS BUFR Clear Radiances
		202	GTS BUFR Retrieved Profiles/Clear Radiances
		210	ATOVS/GRAD
		211	RTOVS
		212	TOVS
		215	SSMI
9	SCATTEROMETER	122	ERS-1, ERS-2
		210	NSCAT
		301	QuikSCAT
		139	ASCAT
10	RAW RADIANCE		

Table 6.2 Some variables (VARNO) in the ODB. Check the ODB governance database <http://apps.ecmwf.int/odbgov> for the latest information.

Code	Name	Unit
3	Wind Component (u)	ms^{-1}
4	Wind Component (v)	ms^{-1}
1	Geopotential (Z)	m^2s^{-2}
57	Thickness (DZ)	m^2s^{-2}
29	Relative Humidity (RH)	numeric
9	Precipitable Water Content (PWC)	kgm^{-2}
58	2 m Relative Humidity (RH_{2m})	numeric
2	Temperature	K
59	Dew Point	K
39	2 m Temperature (T_{2m})	K
40	2 m Dew Point (Td_{2m})	K
11	Surface Temperature (T_s)	K
30	Pressure Tendency (P_t)	Pa/3h
60	Past Weather (W)	WMO Code 4561
61	Present Weather (WW)	WMO Code 4677
62	Visibility (V)	WMO Code 4300
63	Type of High Clouds (C_H)	WMO Code 0509
64	Type of Middle Clouds (C_M)	WMO Code 0515
65	Type of Low Clouds (C_L)	WMO Code 0513
66	Cloud Base Height (N_h)	m
67	Low Cloud Amount (N)	WMO Code 2700
68	Additional Cloud Group Height ($h_s h_s$)	m
69	Additional Cloud Group Type (C)	WMO Code 0500
70	Additional Cloud Group Amount (N_s)	WMO Code 2700
71	Snow Depth (Sd)	m
72	State of Ground (E)	WMO Code 0901
73	Ground Temperature ($T_g T_g$)	K
74	Special Phenomena ($S_p S_p$)	WMO Code 3778
75	Special Phenomena ($s_p s_p$)	WMO Code 3778
76	Ice Code Type (R_s)	WMO Code 3551
77	Ice Thickness ($E_s E_s$)	WMO Code 1751
78	Ice (I_s)	WMO Code 1751
79	Time Period of Rain Information ($t_r t_r$)	hour
80	6 Hour Rain Amount	kgm^{-2}
81	Maximum Temperature (JJ)	K
82	Ship Speed (V_s)	ms^{-1}
83	Ship Direction (D_s)	degree
84	Wave Height ($H_w H_w$)	m
85	Wave Period ($P_w P_w$)	s
86	Wave Direction ($D_w D_w$)	degree
87	General Cloud Group	WMO Code

6.3 VARIABLE CODES (VARNO)

To provide easy recognition of ‘observed’ variables each of them is assigned a numerical code. These codes are then embedded in ODB reports. For illustrative purposes some historical codes are listed in Table 6.2 but for the latest information please check the subroutine or ideally the ODB governance database <http://apps.ecmwf.int/odbgov>. The **VARNO_MODULE** encodes these varnos in the IFS and is auto-generated from the ODB database.

Table 6.2 Continued.

Code	Name	Unit
88	Relative Humidity from Low Clouds	numeric
89	Relative Humidity from Middle Clouds	numeric
90	Relative Humidity from High Clouds	numeric
91	Total Amount of Clouds	WMO Code 20011
92	6 Hour Snowfall	m
110	Surface Pressure (P_s)	Pa
111	Wind Direction	degree
112	Wind Force	ms^{-1}
119	Brightness Temperature (T_b)	K
120	Raw Radiance	K
121	Cloud Amount from Satellite	%
122	Backscatter (σ^0)	dB
5	Wind Shear ($\partial u / \partial z$)	s^{-1}
6	Wind Shear ($\partial v / \partial z$)	s^{-1}
41	$u_{10\text{m}}$	ms^{-1}
42	$v_{10\text{m}}$	ms^{-1}
19	Layer Relative Humidity	numeric
200	Auxiliary Variable	numeric
123	Cloud Liquid Water (Q_l)	kgkg^{-1}
124	Ambiguous v	ms^{-1}
125	Ambiguous u	ms^{-1}
7	Specific Humidity (Q)	kgkg^{-1}
126	Ambiguous Wind Direction	degree
127	Ambiguous Wind Speed	ms^{-1}
8	Vertical Speed	ms^{-1}
56	Virtual Temperature (T_v)	K
206	Ozone	kgm^{-2}
156	Height	m
215	SSM/I Pseudo Variable	kgm^{-2}
160	Past Weather	numeric
130	Pressure Tendency Characteristics	numeric
12	Sea Water Temperature	K
192	Radar Reflectivity	Db
128	Atmospheric Path Delay in Satellite Signal	m
162	Radio Occultation Bending Angle	Rad
187	Horizontal line-of-sight wind component	ms^{-1}
174	Aerosol optical depth at $0.55\mu\text{m}$ (AOD)	
163	Limb Radiances	
181	GEMS reactive gases, N02	
182	GEMS reactive gases, S02	
183	GEMS reactive gases, CO	
184	GEMS reactive gases	
185	GEMS reactive gases, G03	
175	Cloud optical depth (COD)	
176	Ratio of fine mode to total aerosol optical depth at $0.55\mu\text{m}$ (RAO)	
177	Aerosol reflectance multi-channel (RFA)	
178	Aerosol optical depth multi-channel (ODA)	
179	Normalized Soil Moisture	0-100%
180	Soil Moisture	$\text{kg}^3\text{kg}^{-3}$
186	GHG	
187	GHG	
195	Radar doppler wind	

6.4 ADDITIONAL OBSERVATION OPERATOR CODES: NVAR AND CVAR_NAME

NVAR is an index into the Jo table and into VarQC settings: most VARNOs map onto an NVAR. Only if there is a mapping from VARNO to NVAR can an observation be processed in 4D-Var. Previously NVAR was a static and hardcoded set of indices, prone to terrible merge problems. It is now dynamically allocated at runtime, but the resulting NVAR numbering will be static as long as this code does not change. The complete and up to date set of NVAR variable numbers is automatically generated and printed in the IFS 4D-Var log file output.

NVAR is used to index the Jo-table and the VarQC settings (including Huber); at Meteo-France it is the numbering scheme used for the NOTVAR mechanism to remove observations from the processing (a sort of instant blacklist, see hop.F90)

The NVAR-associated character code is used in only two places: (i) in the Jo table printouts; (ii) to select the appropriate obs operator in ppnew.F90 and similar conventional operators.

To add a new one, first add the NVAR definition of your choice. Then add a call to register_nvar

Why do we not use VARNO directly?

- There is no 1-to-1 mapping; for example U and V wind components both map onto NVAR%U
- NVAR is a more compact range of numbers and reduces the size of arrays in Jo

There is a strong case for removing NVAR altogether (and using VARNOs directly), but there are the following issues:

- Meteo France use the NOTVAR function routinely in operations, assuming a fixed set of numerical values (i.e. hardcoded) for each NVAR. Any replacement needs significant consultation or adaptation.
- Improvements are needed in setting up the Huber norm (see defrun.F90) which is indexed on NVAR.
- A transparent alternative way of collapsing U and V onto one Jo variable is needed.
- Use of the NVAR direct pointers (e.g. NVAR%V10) needs to be eliminated

The conventional observation operator uses a short name (CVAR_NAME, three characters), that is indexed to the NVAR. These codes are also used in the diagnostic Jo-table printouts. Some of the traditional operator names are: U, U10, DD, FF, H, H2, T, Z, DZ, LH, T2, TS, RAD, SN, RR, PS, CC, CLW, Q, FFF, S0, X, PWC, TO3 and TCW, numbered sequentially in NVAR. Each number can be referenced by variables such as NVAR%U(= 1), NVAR%U10(= 2) and so on. See table 6.3.

Table 6.3 An illustrative example of the association between NVAR variable numbering and observation operator routines for some of the longstanding observation types. The CVAR-NAMEs also appear in the printed J_0 -table in the log-file. The complete and up to date set of NVAR variable numbers is automatically generated and printed in the IFS log file output.

NVAR	CVAR_NAME	Observation operator routine	Description
1	U	PPUV	Upper air wind components
2	U10	PPUV10M	10-metre wind components
3	DD		Wind direction
4	FF	PPUV	Wind speed
5	H	PPRH	Relative humidity
6	H2	PPRH2M	2-metre relative humidity
7	T	PPT	Temperature
8	Z	PPGEOP	Geopotential
9	DZ	PPGEOP	Thickness
10	LH	PPRH	Layer mean RH (M-France)
11	T2	PPT2M	2-metre temperature
12	TS		Surface temperature (M-France)
13	RAD	RADTR/RADTR_ML	Radiance data
14	SN		Snow (M-France)
15	RR		Rain rate (M-France)
16	PS	PPPS	Surface pressure
17	CC	PPTCC	Cloud cover
18	CLW	PPCLW	Cloud liquid water
19	Q	PPQ	Specific humidity
20	FFF	PPUV10M	10-metre wind speed
21	S0		Sigma_0
22	X		Reserved
23	PWC	PPPWC	Layer water content or TCWV
24	TO3	PPPWC	Layer ozone content
25	TCW		Layer cloud water content
26	RFL	REFLSIM	Radar reflectivity
27	APD	GPSZEN_DELAY	GPS total zenith delay
28	RO	GPSRO_OP	GPS radio occultation
29	HLS	PP_UV(+conversion to HLOS in HOP)	Horizontal line-of-sight winds
30	AOD	AOD_OP	Aerosol optical depth
31	LRA	RTL_HOP_1D	Limb sound radiance

Table 6.4 *SYNOP instrument specification.*

Type	Bit Position	No. of Bits	Value – Description
Instrument Specification	0	10	32 – SYNOP Instrument Code Type
Not Defined	10–30	21	Reserved

Table 6.5 *AIREP instrument specification.*

Type	Bit Position	No. of Bits	Value – Description
Instrument Specification	0	10	23 – AIREP Instrument Code Type
Flight Information	10	4	BUFR Code Table 8004 – Flight Phase
Not Defined	10–30	21	Reserved

Table 6.6 *SATOB instrument specification.*

Type	Bit Position	No. of Bits	Value – Description
Instrument Specification	0	10	60 - GOES 62 – METEOSAT 63 – Indian SATOB 68 – Japan
I1 (Country Name)	10	4	0 – Europe 1 – Japan 2 – USA 3 – USSR 4 – India
I2I2 (Satellite Indicator Figure)	14	8	4 – METEOSAT 177 – Pretoria 0 – GEOS 3 – Japan 20 – India
Not Defined	22–30	8	Reserved

Table 6.7 *DRIBU instrument specification.*

Type	Bit Position	No. of Bits	Value – Description
Instrument Specification	0	10	Not Defined
K1	10	4	Not Defined
K2	14	4	Not Defined
K3	18	4	Not Defined
Not Defined	22–30	8	Reserved

6.5 OBSERVATION CHARACTERISTICS: INSTRUMENT SPECIFICATION AND RETRIEVAL TYPE

Where applicable, [Tables 6.4](#) to [6.10](#) describe in details how the ODB's instrument specification word is structured. Tables provided are for different observation types.

In [Table 6.11](#) the ODB's header retrieval word codes are described.

Table 6.8 TEMP instrument specification.

Type	Bit Position	No. of Bits	Value – Description
Instrument Specification	0	10	Not Defined
Not Defined	10–30	21	Reserved

Table 6.9 PILOT instrument specification.

Type	Bit Position	No. of Bits	Value – Description
Instrument Specification	0	10	Not Defined
A4	10	4	Not Defined
Not Defined	14–30	17	Reserved

Table 6.10 SATEM instrument specification.

Type	Bit Position	No. of Bits	Value – Description
Instrument Specification	0	23	77 777 777B
I3	24	4	WMO Manual On Codes, vol II, section II-4-E-8
I4	28	4	Data processing technique. WMO Manual On Codes, vol II, section II-4-E-9
I2I2	32	7	Satellite name. WMO Manual on Codes, vol II, section II-4-E-7
I1	39	4	Country operating satellite. WMO code 1761
IS	43	7	Instrument specification code. Research Manual 5, Table 7.5
Not Defined	50	18	Reserved

Table 6.11 Satellite retrieval codes.

Retrieval Codes	Description
1	Clear
2	Partly Clear
3	Cloudy

Table 6.12 *Vertical coordinate.*

Vertical Coordinate Codes	Description
1	Pressure (Pa)
2	Height (GPM)
3	Satellite Channel (numeric)
4	Scatterometer Channel (numeric)

Table 6.13 *Pressure codes.*

Pressure Codes	Description
0	Sea Level
1	Station Level
2	850 hPa Geopotential
3	700 hPa Geopotential
4	500 hPa Geopotential
5	1000 GPM Pressure
6	2000 GPM Pressure
7	3000 GPM Pressure
8	4000 GPM Pressure
9	900 hPa Geopotential
10	1000 hPa Geopotential
11	500 hPa Geopotential
12	925 hPa Geopotential

Table 6.14 *Level ID.*

Bit Position	No. of Bits	Value – Description
0	1	1 – Max Wind Level
1	1	1 – Tropopause
2	1	1 – D Part
3	1	1 – C Part
4	1	1 – B Part
5	1	1 – A Part
6	1	1 – Surface Level
7	1	1 – Significant Wind Level
8	1	1 – Significant Temperature Level
9–31	24	Not Defined

6.6 VERTICAL COORDINATE: PRESSURE, SATELLITE ID AND LEVEL ID CODES

In the ODB the vertical coordinate is expressed by various codes, and [Table 6.12](#) describes those codes.

Also, the ODB pressure code word is expressed in terms of codes which are defined in [Table 6.13](#).

Upper air observations (TEMP and PILOT) have the level at which the observation was taken defined in terms what it is and that information is stored in the ODB. Details are given in [Table 6.14](#).

Table 6.15 Report Status.

Bit Position	No. of Bits	Value – Description
0	1	1 – Report Active
1	1	1 – Passive Report
2	1	1 – Rejected Report
3	1	1 – Blocklisted Report

Table 6.16 Blocklist Events.

Bit Position	No. of Bits	Value – Description
0	1	1 – Monthly Monitoring
1	1	1 – Constant Blocklisting
2	1	1 – Experimental Blocklisting
3	1	1 – Whitelisting
4	1	1 – Experimental Whitelisting
5	1	1 – Observation Type Blocklisting
6	1	1 – Station ID Blocklisted
7	1	1 – Code Type Blocklisted
8	1	1 – Instrument Type Blocklisted
9	1	1 – Date Blocklisted
10	1	1 – Time Blocklisted
11	1	1 – Latitude Blocklisted
12	1	1 – Longitude Blocklisted
13	1	1 – Station Altitude Blocklisted
14	1	1 – Blocklisted due to Land/Sea Mask
15	1	1 – Blocklisted due to Model Orography
16	1	1 – Blocklisted due to distance from reference point
17–30	14	Not Used

Table 6.17 Global report events.

Bit Position	No. of Bits	Description (Value)
0	1	1 – No Data in Report
1	1	1 – All Data Rejected
2	1	1 – Bad Reporting Practice
3	1	1 – Rejected due to RDB Flag
4	1	1 – Redundant Report
5	1	1 – Missing Station Altitude
6	1	1 – Failed Quality Control
7	1	1 – Report Overcast IR (deprecated)

6.7 ODB REPORT STATUS: EVENTS, FLAGS AND CODES

The status of each ODB report is described in terms of being active, passive, rejected or blocklisted. For some microwave radiances, the additional flag `use_emiskf_only` is also used, see 3.3.2 and (b). The ODB report status word is packed with the 4 bits given in Table 6.15.

There is one, 31 bits packed, word for each ODB report to account for various blocklist events. Details are given in Table 6.16.

Each ODB report has two words to store report events. Each report event word uses 31 bits. These events are set during observation processing to describe in more details what happened with a report.

The first ODB report event word is described in Table 6.17.

Table 6.18 *TEMP report events.*

Bit Position	No. of Bits	Value – Description
0	1	1 - Old Style Z Bias Correction Applied
1	1	1 - New Style T Bias Correction Applied
2	1	1 - RH Bias Correction Applied
3–30	28	Not Used

Table 6.19 *PILOT report events.*

Bit Position	No. of Bits	Value – Description
0	1	1 - American Wind Profiler
1	1	1 - European Wind Profiler
2–30	29	Not Used

Table 6.20 *SATEM report events.*

Bit Position	No. of Bits	Value – Description
0	1	1 - Thinned Report
1–30	30	Not Used

Table 6.21 *SCAT report events.*

Bit Position	No. of Bits	Value – Description
0	1	1 - Report thinned in across-node direction
1	1	1 - Reported Wind Directions too Close
2	1	1 - Report in QuikScat outer swath
3	1	1 - Report Contaminated by Rain
4–30	29	Not Used

The second ODB report event word holds an additional set of events which are now dependent on observation type. Details are given in [Tables 6.18](#) to [6.21](#).

The ODB report RDB flag word ([DEPRECATED], probably unused) is 30 bits packed which contains flags for five report parameters: latitude, longitude, date, time and altitude. Each parameter occupies 6 bits with further stratification which is identical for every parameter as indicated in [Table 6.22](#).

Table 6.22 RDB report (latitude, longitude, date, time and altitude) flags. [DEPRECATED], probably unused.

Parameter	No. of Bits	Bit Position	Bit Position	No. of Bits	Value – Description
			0	1	0 – No Human Monitoring Substitution 1 – Human Monitoring Substitution
Latitude	6	0+	+1	1	0 – No Q/C Substitution 1 – Q/C Substitution
Longitude	6	6+			
Date	6	12+	+2	1	0 – Override Flag not Set 1 – Override Flag Set
Time	6	18+			
Altitude	6	24+	+3	2	0 – Parameter Correct 1 – Parameter Probably Correct 2 – Parameter Probably Incorrect 3 – Parameter Incorrect
			+5	1	0 – Parameter Flag Set by Q/C or not Checked 1 – Parameter Flag Set by Human Monitoring

Table 6.23 *Datum status.*

Bit Position	No. of Bits	Value – Description
0	1	1 – Report Active
1	1	1 – Passive Report
2	1	1 – Rejected Report
3	1	1 – Blocklisted Report

Table 6.24 *Global datum events.*

Bit Position	No. of Bits	Value – Description
0	1	1 – Missing Vertical Coordinate
1	1	1 – Missing Observed Value
2	1	1 – Missing Background (First Guess) Value
3	1	1 – Rejected due to RDB Flag
4	1	1 – Activated due to RDB Flag
5	1	1 – Activated by Whitelist
6	1	1 – Bad Reporting Practice
7	1	1 – Vertical Position out of Range
8	1	1 – Reference Level Position out of Range
9	1	1 – Too Big First Guess Departure
10	1	1 – Too Big Departure in Assimilation
11	1	1 – Too Big Observation Error
12	1	1 – Redundant Datum
13	1	1 – Redundant Level
14	1	1 – Report Over Land
15	1	1 – Report Over Sea
16	1	1 – Not Analysis Variable
17	1	1 – Duplicate Datum/Level
18	1	1 – Too Many Surface Data
19	1	1 – Multi Level Check
20	1	1 – Level Selection
21	1	1 – Vertical Consistency Check
22	1	1 – Vertical Coordinate Changed from Z to P
23	1	1 – Datum Rejected via Namelist
24	1	1 – Combined Flagging
25	1	1 – Datum Rejected due to Rejected Report
26	1	1 – Variational QC Performed
27	1	1 – Observation Error Increased
28	1	1 – Cloud Contamination
29	1	1 – Rain Contamination
30	1	1 – Aerosol Contamination
31	1	1 – Missing or Not Sensible Emissivity Values

6.8 DATUM STATUS: EVENTS, RDB AND ANALYSIS FLAGS

The status of each datum, like report status, is described in terms of being: active, passive, rejected or blocklisted. [Table 6.23](#) shows that the ODB datum status is a packed word with 4 bits used to describe its status.

There are two ODB words reserved for datum events. They both use 31 bits each to store relevant information. The first event word has the same structure for all observation types, whereas the second event word is observation type dependent. [Tables 6.24](#) to [6.27](#) describe the event words structures for the observation types that use them.

Table 6.25 *SYNOP datum events.*

Bit Position	No. of Bits	Value – Description
0	1	1 – Bias Corrected Ps
1–30	30	Not Used

Table 6.26 *DRIBU datum events.*

Bit Position	No. of Bits	Value – Description
0	1	1 – Bias Corrected Ps
1–30	30	Not Used

Table 6.27 *TEMP datum events.*

Bit Position	No. of Bits	Value – Description
0	1	1 – Bias Corrected Value Used
1–30	30	Not Used

Table 6.28 *Datum blocklist events.*

Bit Position	No. of Bits	Value – Description
0	1	1 – Pressure Blocklisted
1	1	1 – Variable Blocklisted
2	1	1 – Blocklisted due to Pressure Code
3	1	1 – Blocklisted due to Distance from Reference Point
4	1	1 – Blocklisted due to Type of Vertical Coordinate
5	1	1 – Blocklisted due to Observed Value
6	1	1 – Blocklisted due to First Guess departure
7–30	24	Not Used

Furthermore, each datum in the ODB has a blocklist event word. This word uses 31 bits to describe various blocklist events as indicated in [Table 6.28](#).

For each datum in ODB there is an RDB flag ([DEPRECATED], probably unused) word which holds flags for pressure (vertical coordinate) and the datum itself. This is packed word with 30 bits used – see [Table 6.29](#). Pressure and datum RDB flags use 15 bits each. Thus pressure RDB flag starts at bit position 0, whereas the datum flag starts at bit position 15. Each 15 bits structure is further stratified in exactly the same way for both parameters:

In addition to RDB datum flags there is a word in ODB to store analysis flags. There are five types of analysis flags: final analysis, first guess, departure, variational q/c and blocklist flags. Each flag occupies 4 bits and the exact description is given in [Table 6.30](#).

Table 6.29 RDB pressure (vertical coordinate) and datum flags. [DEPRECATED], probably unused.

Parameter	No. of Bits	Bit Position	Bit Position	No. of Bits	Value – Description
			0	1	0 – No Human Monitoring Substitution 1 – Human Monitoring Substitution
Pressure	15	0+	+1	1	0 – No Q/C Substitution 1 – Q/C Substitution
Datum	15	15+	+2	1	0 – Override Flag not Set 1 – Override Flag Set
			+3	2	0 – Correct 1 – Probably Correct 2 – Probably Incorrect 3 – Parameter Incorrect
			+5	1	0 – Flag Set by Q/C or not Checked 1 – Flag Set by Human Monitoring
			+6	2	0 – Previous Analysis judged it correct 1 – Previous Analysis judged it probably correct 2 – Previous Analysis judged it probably incorrect 3 – Previous Analysis judged it incorrect
			+8	1	0 – Not used by previous analysis 1 – Used by previous analysis
			+9	5	Not Used

Table 6.30 *Analysis flags. DEPRECATED, probably unused.*

Flag Type	Bit Position	No. of Bits	Value – Description
Final	0	4	0 – Correct 1 – Probably correct 2 – Probably incorrect 3 – Incorrect
First Guess	4	4	0 – Correct 1 – Probably correct 2 – Probably incorrect 3 – Incorrect
Departure	8	4	0 – Correct 1 – Probably correct 2 – Probably incorrect 3 – Incorrect
Variational Q/C	12	4	0 – Correct 1 – Probably correct 2 – Probably incorrect 3 – Incorrect
Blocklist	16	4	0 – Correct 1 – Probably correct 2 – Probably incorrect 3 – Incorrect
Not Defined	20	11	Reserved

References

- Agusti-Panareda, A., Vasiljevic, D., Beljaars, A., Bock, O., Guichard, F., Nuret, M., Mendez, A. G., Andersson, E., Bechtold, P., Fink, A., Hersbach, H. *et al.* (2009). Radiosonde humidity bias correction over the west african region for the special amma reanalysis at ecmwf. *Quarterly Journal of the Royal Meteorological Society*, doi:10.1002/qj.396.
- Andersson, E., Pailleux, J., Thépaut, J.-N., Eyre, J. R., McNally, A. P., Kelly, G. A. and Courtier, P. (1994). Use of cloud-cleared radiances in three/four-dimensional variational data assimilation. *Q. J. R. Meteorol. Soc.*, **120**, 627–653.
- Baordo, F. and Geer, A. J. (2016). Assimilation of SSMIS humidity-sounding channels in all-sky conditions over land using a dynamic emissivity retrieval. *Quarterly Journal of the Royal Meteorological Society*, doi:10.1002/qj.2873.
- Bauer, P., Geer, A. J., Lopez, P. and Salmond, D. (2010). Direct 4D-Var assimilation of all-sky radiances: Part I. Implementation. *Q. J. R. Meteorol. Soc.*, **136**, 1868–1885.
- Bonavita, M., Dahoui, M., Lopez, P., Prates, F., Holm, E., Chiara, G. D., Geer, A., Isaksen, L. and Ingleby, B. (2017). On the initialization of tropical cyclones. *ECMWF Tech. Memo. No. 810*.
- Bormann, N. (2014). Reformulation of the emissivity kalman filter atlas. *RD Internal Memorandum*, p. 14pp, available online: <http://intra.ecmwf.int/publications/library/do/references/show?id=1471>.
- Bormann, N. (2017). Slant path radiative transfer for the assimilation of sounder radiances. *Tellus A: Dynamic Meteorology and Oceanography*, **69**(1), 1272779.
- Bormann, N. and Bonavita, M. (2013). Spread of the ensemble of data assimilations in radiance space. *ECMWF Technical Memorandum 708*.
- Bormann, N., Bonavita, M., Dragani, R., Eresmaa, R., Matricardi, M. and McNally, A. (2016). Enhancing the impact of IASI observations through an updated observation-error covariance matrix. *Q. J. R. Meteorol. Soc.*, **142**, 1767–1780.
- Bormann, N. and Healy, S. (2006). A fast radiative-transfer model for the assimilation of MIPAS limb radiances: Accounting for horizontal gradients. *Q. J. R. Meteorol. Soc.*, **132**, 2357–2376.
- Bormann, N., Healy, S. and Hamrud, M. (2007). Assimilation of MIPAS limb radiances in the ECMWF system. Part II: Experiments with a 2-dimensional observation operator and comparison to retrieval assimilation. *Q. J. R. Meteorol. Soc.*, **133**, 329–346.
- Bormann, N., Matricardi, M. and Healy, S. (2005). A fast radiative transfer model for the assimilation of infrared limb radiances from MIPAS. *Q. J. R. Meteorol. Soc.*, **131**, 1631–1653.
- Bormann, N., Saarinen, S., Kelly, G. and Thépaut, J.-N. (2003). The spatial structure of observation errors in atmospheric motion vectors from geostationary satellite data. *Mon. Wea. Rev.*, **131**, 706–718.
- Bormann, N. and Thépaut, J.-N. (2007). Assimilation of MIPAS limb radiances in the ECMWF system. Part I: Experiments with a 1-dimensional observation operator. *Q. J. R. Meteorol. Soc.*, **133**, 309–327.
- Burrows, C. and C.Lupu (2024). A new way of handling observation errors for infrared measurements. *ECMWF Newsletter No. 180*, pp. 9–10.

- Cardinali, C., Isaksen, L. and Andersson, E. (2003). Use and impact of automated aircraft data in a global 4D-Var data assimilation system. *Mon. Wea. Rev.*, **131**, 1865–1877.
- Courtier, P., Andersson, E., Heckley, W., Pailleux, J., Vasiljevic, D., Hamrud, M., Hollingsworth, A., Rabier, F. and Fisher, M. (1998). The ECMWF implementation of three dimensional variational assimilation (3D-Var). I: Formulation. *Q. J. R. Meteorol. Soc.*, **124**, 1783–1807.
- Duncan, D. I., Bormann, N., Geer, A. J. and Weston, P. (2022). Assimilation of AMSU-A in all-sky conditions. *Mon. Weather Rev.*, **150**(5), 1023 – 1041, doi:10.1175/MWR-D-21-0273.1.
- Duncan, D. I., Bormann, N., Geer, A. J. and Weston, P. (2024). Superobbing and thinning scales for all-sky humidity sounder assimilation. *Mon. Weather Rev.*, **152**(8), 1821 – 1837, doi:10.1175/MWR-D-24-0020.1.
- Engelen, R. J., Andersson, E., Chevallier, F., Hollingsworth, A., Matricardi, M., McNally, A., Thépaut, J.-N. and Watts, P. (2004). Estimating atmospheric CO₂ from advanced infrared satellite radiances within an operational 4D-Var data assimilation system: Methodology and first results. *J. Geophys. Res.*, **109**, D19309.
- Eyre, J. R. (1991). A fast radiative transfer model for satellite sounding systems. *ECMWF Tech. Memo. No. 176*.
- Fielding, M. and Stiller, O. (2019). Characterizing the representativity error of cloud profiling observations for data assimilation. *J. Geophys. Res. - Atmospheres*, **124**, 4086–4103.
- Geer, A. J. (2024a). Joint estimation of sea ice and atmospheric state from microwave imagers in operational weather forecasting. *Quarterly Journal of the Royal Meteorological Society*, **150**(763), 3796–3826.
- Geer, A. J. (2024b). Simultaneous inference of sea ice state and surface emissivity model using machine learning and data assimilation. *Journal of Advances in Modeling Earth Systems*, **16**(7), e2023MS004080.
- Geer, A. J., Baordo, F., Bormann, N. and English, S. (2014). All-sky assimilation of microwave humidity sounders. *ECMWF Tech. Memo. No. 741*.
- Geer, A. J. and Bauer, P. (2011). Observation errors in all-sky data assimilation. *Q. J. R. Meteorol. Soc.*, DOI: 10.1002/qj.830.
- Geer, A. J., Bauer, P. and Lopez, P. (2008). Lessons learnt from the operational 1D+4D-Var assimilation of rain- and cloud-affected SSM/I observations at ECMWF. *Q. J. R. Meteorol. Soc.*, **134**, 1513–1525.
- Geer, A. J., Bauer, P. and Lopez, P. (2010). Direct 4D-Var assimilation of all-sky radiances: Part II. Assessment. *Q. J. R. Meteorol. Soc.*, **136**, 1886–1905.
- Geer, A. J., Lonitz, K., Duncan, D. and Bormann, N. (2022). Improved surface treatment for all-sky microwave observations. *ECMWF Tech. Memo. No. 894*.
- Geleyn, J.-F. (1988). Interpolation of wind, temperature and humidity values from the model levels to the height of measurement. *Tellus*, **40**, 347–351.
- Han, W. and Bormann, N. (2016). Constrained adaptive bias correction for satellite radiance assimilation in the ECMWF 4D-Var system. *ECMWF Tech. Memo. No. 783*, doi:10.21957/rex0omex.
- Healy, S. and Thépaut, J.-N. (2006). Assimilation experiments with CHAMP GPS radio occultation measurements. *Q. J. R. Meteorol. Soc.*, **132**, 605–623.
- Hersbach, H. (2010a). Assimilation of scatterometer data as equivalent-neutral wind. *ECMWF Tech. Memo. No. 629*.
- Hersbach, H. (2010b). Comparison of c-band scatterometer cm05.n equivalent neutral winds with ecmwf. *J. Atmos. Oceanic Technol.*, **27**, 721–736.
- Hersbach, H., Stoffelen, A. and de, S. H. (2007). An improved C-band scatterometer ocean geophysical model function: CMOD5. *J. Geophys. Res.*, **112**, doi:10.1029/2006JC003743.

- Hocking, J. (2014). Interpolation methods in the rrtov radiative transfer model. *Forecast Research Technical Report No. 590*.
- Huddleston, J. N. and Stiles, B. W. (2000). Multidimensional histogram (MUDH) rain flag. *Product Description Version 2.1*, URL <http://podaac-www.jpl.nasa.gov/quikscat/>.
- Ingleby, B. (2015). Global assimilation of air temperature, humidity, wind and pressure from surface stations. *Q. J. R. Meteorol. Soc.*, **141**, 504–517.
- Ingleby, B. (2017). An assessment of different radiosonde types 2015/2016. *ECMWF Tech. Memo. No. 807*.
- Ingleby, B. (2025). The benefits and challenges of using dense mode-s aircraft winds at ecmwf. *Quarterly Journal of the Royal Meteorological Society*, URL <http://doi.org/10.1002/qj.70032>.
- Ingleby, B., Arduini, G., Balsamo, G., Bousetta, S., Ochi, K., Pinnington, E. and de Rosnay, P. (2024). Improved two-metre temperature forecasts in the 2024 upgrade. *ECMWF Newsletter No. 178*, pp. 24–29.
- Isaksen, L. and Janssen, P. A. E. M. (2004). Impact of ERS scatterometer winds in ECMWF's assimilation. *Q. J. R. Meteorol. Soc.*, **130**, 1793–1814.
- Järvinen, H., Andersson, E. and Bouttier, F. (1999). Variational assimilation of time sequences of surface observations with serially correlated errors. *ECMWF Tech. Memo. No. 266*.
- Järvinen, H., Saarinen, S. and Undén, P. (1996). *User's Guide for Blacklisting*. Available on request from ECMWF, Shinfield Park, RG2 9AX, Reading, Berkshire, UK.
- Järvinen, H. and Undén, P. (1997). Observation screening and first guess quality control in the ECMWF 3D-Var data assimilation system. *ECMWF Tech. Memo. No. 236*.
- Karbou, F., Gérard, E. and Rabier, F. (2006). Microwave land emissivity and skin temperature for AMSU-A and -B assimilation over land. *Q. J. R. Meteorol. Soc.*, **132**, 2333–2355.
- Kilic, L., Prigent, C., Jimenez, C., Turner, E., Hocking, J., English, S., Meissner, T. and Dinnat, E. (2023). Development of the SURface Fast Emissivity Model for Ocean (SURFEM-Ocean) based on the PARMIO radiative transfer model. *Earth Space Sci.*, **10**(11), e2022EA002785, doi:10.1029/2022EA002785.
- Krzeminski, B., Bormann, N., Karbou, F. and Bauer, P. (2009a). Improved use of surface-sensitive microwave radiances at ECMWF. In *Proceedings of 2009 EUMETSAT Meteorological Satellite Conference*, Bath, United Kingdom.
- Krzeminski, B., Bormann, N., Kelly, G., McNally, A. and Bauer, P. (2009b). Revision of the HIRS cloud detection at ecmwf. *ECMWF/EUMETSAT fellowship reports No. 19*.
- Lean, K. and Bormann, N. (2021). Using model cloud information to reassign low level AMVs for NWP. *ECMWF/EUMETSAT fellowship reports No. 56*.
- Lean, P., Geer, A., Salmond, D. and Hague, J. (2016). The cost and scalability of observation processing in IFS; implications for future observation usage. *RD internal memo, RD16-045*.
- Leidner, S. M., Hoffman, R. N. and Augenbaum, J. (2000). *SeaWinds Scatterometer Real-Time BUFR Geophysical Data Product, User's Guide Version 2.3.0*. NOAA/NESDIS.
- Lonitz, K. and Geer, A. (2015). New screening of cold-air outbreak regions used in 4D-Var all-sky assimilation. *EUMETSAT/ECMWF Fellowship Programme Research Reports*, (35), URL <https://www.ecmwf.int/node/10777>.
- Lönnerberg, P. (1989). Developments in the ECMWF analysis system. In *Proc. ECMWF Seminar on Data assimilation and the Use of Satellite Data*, pp. 75–119, Reading, 5–9 September 1988.
- Lönnerberg, P. and Shaw, D. (1985). Data selection and quality control in the ECMWF analysis system. In *ECMWF Workshop on The Use And Quality Control of Meteorological Observations*, pp. 225–254, Reading, 6–9 November 1984.

- Lönnerberg, P. and Shaw, D. (Eds) (1987). *ECMWF Data Assimilation Scientific Documentation, Research Manual 1*.
- Lopez, P. (2011). Direct 4d-var assimilation of NCEP stage IV radar and gauge precipitation data at ecmwf. *Monthly Weather Review*, **139**(7), 2098–2116.
- Lorenc, A. C. (1986). Analysis methods for numerical weather prediction. *Q. J. R. Meteorol. Soc.*, **112**, 1177–1194.
- Matricardi, M., Chevallier, F. and Tjemkes, S. (2001). An improved fast radiative transfer model for the assimilation of radiance observations. *ECMWF Tech. Memo. No. 345*.
- McNally, A. P., Andersson, E., Kelly, G. and Saunders, R. W. (1999). The use of raw TOVS/ATOVS radiances in the ECMWF 4D-Var assimilation system. *ECMWF Newsletter No. 83*, pp. 2–7.
- McNally, A. P. and Watts, P. D. (2003). A cloud detection algorithm for high spectral resolution infrared sounders. *Q. J. R. Meteorol. Soc.*, **129**, 3411–3423.
- Munro, R., Köpken, C., Kelly, G., Thépaut, J.-N. and Saunders, R. (2004). Assimilation of Meteosat radiance data within the 4D-Var system at ECMWF: Data quality monitoring, bias correction and single-cycle experiments. *Q. J. R. Meteorol. Soc.*, **130**, 2293–2313.
- Pailleux, J. (1990). A global variational assimilation scheme and its application for using TOVS radiances. In *Proc. WMO International Symposium on Assimilation of Observations in Meteorology and Oceanography*, pp. 325–328, Clermont-Ferrand, France.
- Pauley, P. and Ingleby, B. (2022). *Assimilation of in-situ observations*. Springer, URL <https://link.springer.com/book/10.1007/978-3-030-77722-7>.
- Saunders, R., Hocking, J., Turner, E., Rayer, P., Rundle, D., Brunel, P., Vidot, J., Roquet, P., Matricardi, M., Geer, A., Bormann, N. and Lupu, C. (2018). An update on the RTTOV fast radiative transfer model (currently at version 12). *Geosci. Model Dev.*, **11**, 2717–2737, doi:10.5194/gmd-11-2717-2018.
- Saunders, R. W. and Matricardi, M. (1998). A fast forward model for ATOVS (RTATOV). In *Tech. Proc. 9th International TOVS Study Conference*, p. 11 pp, Igls, Austria, 20–26 February, 1997.
- Scanlon, T., Geer, A., Bormann, N. and Browne, P. (2024). Improving ocean surface temperature for nwp using all-sky microwave imager observations. *EUMETSAT/ECMWF Fellowship Programme Research Report*, (RR64), doi:10.21957/c16be07b23.
- Sherlock, V. J. (1999). ISEM-6: Infrared Surface Emissivity Model for RRTOV-6. *Forecasting Research Technical Report No. 299*, available from Met Office, London Road, Bracknell, Berkshire RG12 2SZ, UK.
- Simmons, A. J. and Burridge, D. (1981). An energy and angular momentum conserving vertical finite difference scheme and hybrid coordinate. *Mon. Wea. Rev.*, **109**, 758–766.
- Simmons, A. J. and Chen, J. (1991). The calculation of geopotential and the pressure gradient in the ECMWF atmospheric model: Influence on the simulation of the polar atmosphere and on temperature analyses. *Q. J. R. Meteorol. Soc.*, **117**, 29–58.
- Sonntag, D. (1994). Advancements in the field of hygrometry. *Z. Meteorol.*, **3**(2), 51–66.
- Stoffelen, A. and Anderson, D. (1997). Ambiguity removal and assimilation of scatterometer data. *Q. J. R. Meteorol. Soc.*, **123**, 491–518.
- Tomassini, M., LeMeur, D. and Saunders, R. W. (1998). Near-surface satellite wind observations of hurricanes and their impact on ECMWF model analyses and forecasts. *Mon. Wea. Rev.*, **126**, 1274–1286.
- Vasiljevic, D. (2006). Handling biases in surface pressure (ps) observations in data assimilation. *ECMWF/EUMETSAT NWP-SAF workshop on bias estimation in data assimilation*.

Wattrelot, E., Caumont, O. and Mahfouf, J.-F. (2014). Operational implementation of the 1D+3D-Var assimilation method of radar reflectivity data in the AROME model. *Monthly Weather Review*, **142**(5), 1852–1873.

WMO (2024). *Guide to instruments and methods of observation*, WMO-No. 8. URL https://community.wmo.int/en/activity-areas/imop/wmo-no_8.