Australian Government
Australian Transport Safety Bureau

Methodology for determining the operational reliability of Australian aviation weather forecasts

ATSB Transport Safety Report
Aviation Research Investigation
AR-2013-200 – Accompany document
FINAL – 10 July 2017

Table of Contents

Data sources and software

Data sources

Weather forecasts and reports

All weather forecasts (TAFs[1] and TTFs[2]) and reports (METARs or SPECIs)[3] were provided by the Bureau of Meteorology (BoM) to the ATSB for Mildura and all major Class C airports between 2009 and 2013.

METAR and SPECI data was provided in a raw format to allow an evaluation of any changes that may have occurred in the system and to give the closest reflection of actual conditions. This was in a similar structure to that released by Airservices Australia for use by the aviation community. However, METAR and SPECI data was provided in higher fidelity as measured by aerodrome equipment prior to numerical rounding. These were:

- temperature, dewpoint, sea level pressure to one-tenth of a unit (eg. QNH reported as 1013.3 hPA rather than 1013 hPA)
- cloud types and specific oktas for each cloud layer (eg. 2 oktas Cumulus 3200 feet instead of FEW032)
- All maximum wind gusts, regardless of the magnitude rather than only when the maximum wind speed is at least 10 knots above the mean (eg 21018G22KT rather than 21018KT).

All weather data was received in either a text or pdf file, with each forecast and report starting on a new line. No delimiters were provided (or requested) within each line of data to separate different data elements, such as wind and temperature. Techniques developed to extract this data for analysis are described in ‘Formation of weather flat data files’ on page 9 and ‘Formation of relational weather data framework’ on page 19.

Aircraft specifications

Aircraft specifications were combined with the published instrument approach data to determine the landing and alternate minima that would apply to the selected aircraft. Aircraft models were selected based on those models used by large Australian-based commercial public transport operators.

Although the list of aircraft identified for use in the study was not exhaustive, it was expected to be sufficient for input into this analysis. This was due to most published approaches being divided into only one of four aircraft performance categories, based on landing speed. As such, it was expected that this study would be representative of most high capacity transport category aircraft.

This data was gathered from various sources, including the aircraft manufacturer’s flight manuals, operator’s websites, and reference books. The calculation and extraction of this data is discussed in ‘Compilation of aircraft specification flat data file’ on page 16.

---

[1] A TAF is an aviation forecast for expected weather conditions within 5 nautical miles of an aerodrome

[2] A TTF referred to as a trend forecast contains statements of trend for expected weather conditions. The TTF supersedes a TAF for its validity period (typically 3 hours in Australia). These are released as a type of weather report (METAR or SPECI).

[3] METARs and SPECIs refer to routine and special (non-routine) weather reports released for an aerodrome based on recent observations.

Aerodrome instrument approach procedures

A complete set of instrument approach procedural charts were downloaded from the Airservices Australia website in portable document format (pdf) for every aerodrome of interest in the study. These were sourced within the Departure and Approach Procedures (DAP) section of the publically available Aeronautical Information Package (AIP).

The data contained in these procedures was used to define the landing and alternate minima criteria for reported and forecast weather conditions applying to a specified aircraft model. This was used to assess if aerodrome weather forecasts sufficiently predicted conditions from a safety perspective. The specific benchmark used for this study was where forecasting of conditions was above the alternate minima (not requiring a contingency plan), and observed conditions fell below the landing minima (conditions below the published limits for a safe landing based on aircraft and aerodrome equipment). This is discussed in depth in sections ‘Instrument approach selection for each METAR and SPECI report using landing minima criteria’ on page 58, and ‘Assessment of TAFs and TTFs against alternate minima criteria’ on page 67.

The extraction of instrument approach procedure data is discussed in the section ‘Formation of instrument approach procedure flat data file’ on page 14 and ‘Formation of relational instrument approach procedure data framework’ on page 46.

Runway details

The analysis used airport data routinely provided to the ATSB from Airservices Australia for runway details such as length, width and nominal direction. Runway data used in the study was that current in the Australian Aeronautical Information Package on 21 August 2014. This data was provided in the form of a comma separated variable table.

One data field not contained in this table was the precise magnetic direction of the runway. This was required for alignment of wind direction for elements such as calculating the runway direction and the crosswind and tailwind wind components. For aerodromes in this study, the runway magnetic direction was appended to this table using data in the aerodrome instrument approach procedures noted above.

Waypoint information

Waypoint information regarding the co-ordinates, ICAO code and full location name of the waypoint was imported from the ATSB’s data holdings. This table was used to create a link from the aerodrome to runways, runway approaches and landing and alternate minima criteria. This was used as the ‘highest level’ table in the aerodrome instrument approach procedure database described above.

Civil twilight start and end times

Data for the beginning and end of civil twilight was retrieved from the Geoscience Australia website[4] for each aerodrome in the study. Location data was extracted using the National Gazetteer of Australia for the specific locations of interest. The Geoscience Australia algorithm was then used to compute the times. Times were extracted in a fixed UTC offset for the local standard time, and required conversions to be made to daylight savings time where applicable.

Data was imported in the form of a structured text file requiring re-arrangement to compile all times in a formalised computer time format.

---

[4] At the time of writing, the website link was www.ga.gov.au.

Aircraft arrivals

Counts of aircraft arrivals were provided by Airservices Australia for every location. Data was aggregated for every hour of day (local time) from 2010 to 2013 for all requested airports. Arrivals were also divided by the following weight categories:

- Less than 5.7 tonnes
- 5.7 to 15 tonnes
- 15 to 50 tonnes
- 50 to 100 tonnes
- Over 100 tonnes

The ATSB were advised that data for the year 2009 was not available at the time for the airports requested.

Software packages

Microsoft powershell

Powershell is a task automation and configuration management tool from Microsoft, consisting of a command-line shell and associated scripting language built on the .NET framework.

Powershell was chosen and used due to its ability and availability to perform the tasks on unformatted text efficiently. This was used in the preliminary data shaping of raw text files from the Bureau of Meteorology and the Geoscience Australia website.

Microsoft SQL Server Management Studio

Microsoft SQL Server Management Studio 2012 was used as the primary tool for extraction, simulation and outputting of results. As such, the majority of the methodology chapter will refer to processes created utilising structured query language (SQL).

SQL was selected due to it supporting the formation of a relational data structure and availability.

Microsoft Excel

Microsoft Excel was used to produce most results graphs and format data tables.

Methodology of data extraction, formation and evaluation

The analysis performed used a number of data sources noted above drawn together into a single relational database. Prior to being combined, a number of processes were required to manipulate the data into a useable format. Some of the data required significant structural changes, such as being entered manually into a flat table structure and then formed into a relational data structure, whereas other data sources required minimal processes.

The processes described in the following chapter are presented in the same order in which the finalised process was executed using the computer algorithm. This follows a similar, but not identical order to the formation of the computer algorithm. This was due to the iterative nature of the development process requiring retrospective changes to structure and content of some data tables and analyses. An overview of these processes is shown in Figure 1. The processes listed also correspond to major headings in the following sections.

Figure 1: Overview of methodology

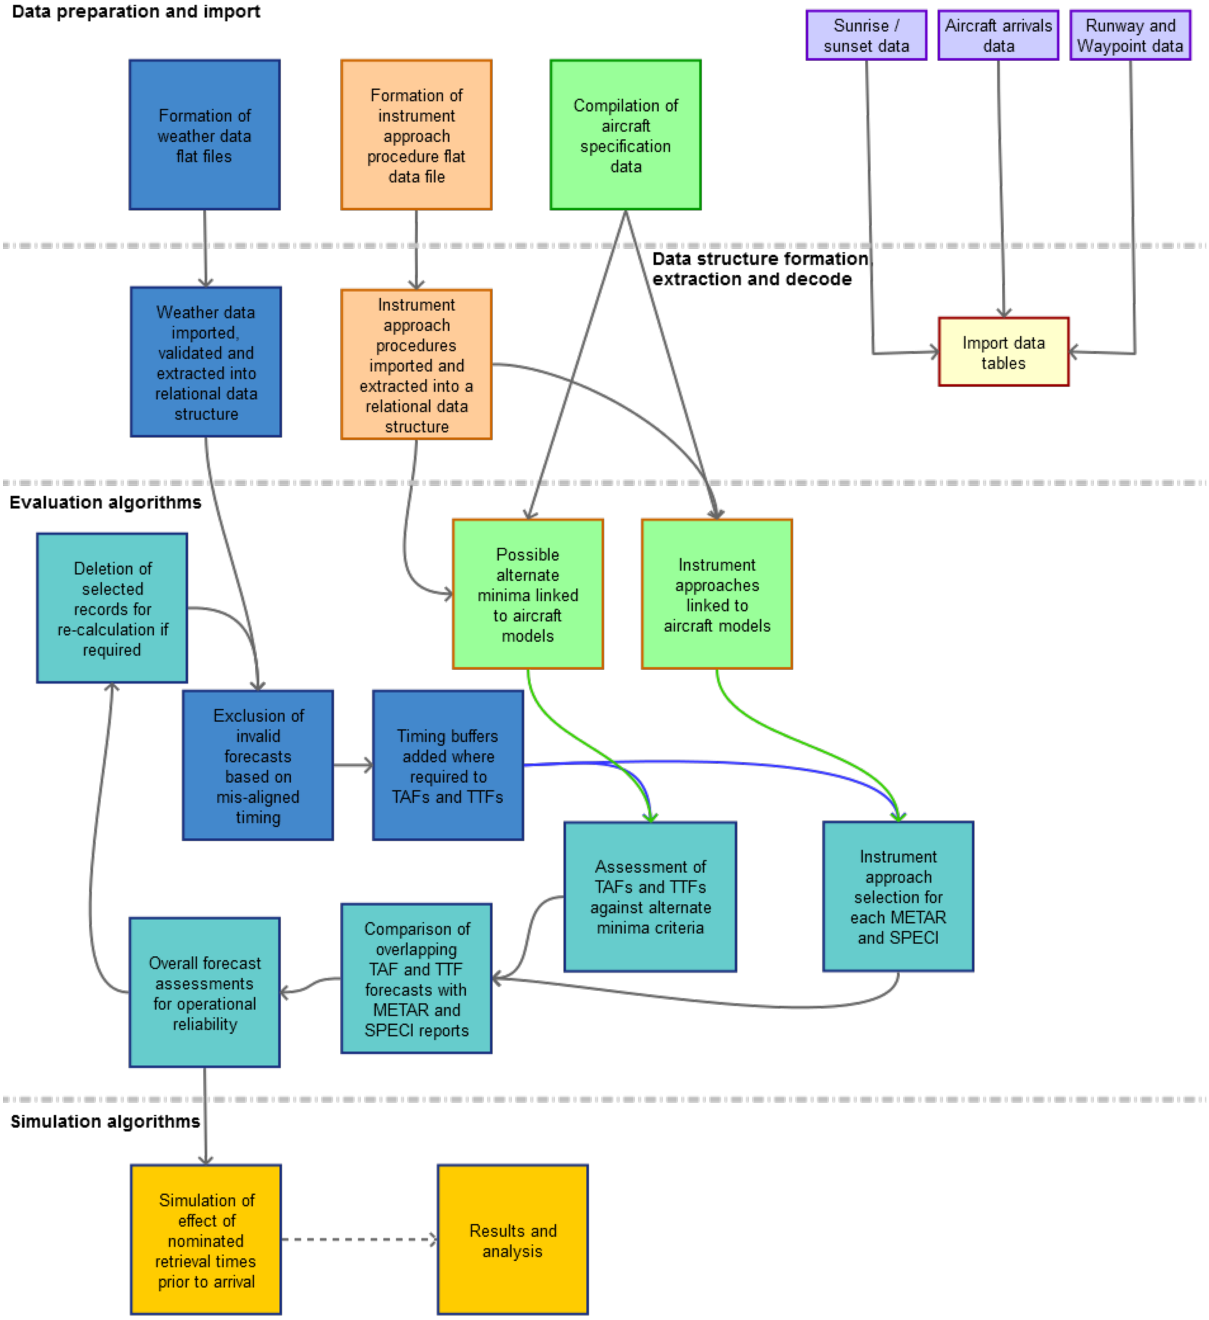

An SQL database was created using a series of stored procedures and functions created specifically for this task. Procedures and functions were used instead of a single script for a number of reasons, primarily:

- to improve transparency of processes where each procedure is designed to perform a small number of tasks
- to provide ease in testing and troubleshooting of errors
- to ensure efficiency in code for repetitive operations, noting that functions and procedures were drawn on multiple times throughout the entire process.

The procedures make use of variables within ‘dynamic sql’ statements to allow for the same procedures to be used between locations, types of aircraft, forecast types and time ranges. The overall program runs three separate SQL scripts, which are in turn used to run each of the main procedures in the extraction, simulation and analysis.

Mechanism to reset records

An SQL procedure was written to remove records from the data import, decode and evaluation stages of the algorithm. At the time of writing, records could be deleted on a location by location basis, allowing records for other locations to be preserved. This enabled development of further analyses techniques for locations yet to be evaluated, while preserving work already done.

Data integrity was maintained using the referential integrity of the database by preventing records from being deleted with ‘child items’ attached. This prevented items from being inadvertently evaluated without having an encoded link back to the source document.

Data preparation and import

Formation of weather flat data files

To enable the comparison of forecasts with weather reports (TAFs and TTFs with METARs and SPECIs), individual elements of each weather report or forecast were required to be extracted from raw text files provided by the Australian Bureau of Meteorology, as listed in ‘Data sources’ above.

The following section shows the processes involved in preparing weather flat data files. This involved delimiting data from the Bureau of Meteorology into two flat data tables (one row per forecast or observation) using Microsoft Powershell. These tables contained:

a. METAR, SPECI and TTF data
b. TAF data.

Delimitation of raw METAR, SPECI and TTF data into a delimited flat file

The following describes the algorithm used to extract and format data from an unformatted but structured text file containing METAR and SPECI reports, which also contain trend forecasts (TTFs) at specific aerodromes. To perform this task Microsoft Powershell was selected for use due to its availability and ability to perform the tasks on the unformatted text efficiently.

To allow manipulation in later processes, data was identified, grouped and delimited into the main elements of these weather products. These groups are shown in Table 1.

Table 1: METAR, SPECI and TTF data flat file column names

Field type	Field name
Unique identifier group	Date and time
Non-standard text	Notes and error text
Unique identifier groups	Report type code
	Report type
	ICAO code
	Issue time
METAR / SPECI Weather groups	Wind
	Visibility
	Runway visual range
	Weather phenomena
	Cloud types and cover
	Temperature and dewpoint
	Sea level pressure (QNH)
	Recent weather phenomena
	Windshear
	Rainfall[5]
	Remarks
Entire forecast	TTF

The final data output table was a tab delimited file with one row per METAR or SPECI report including any attached TTF.

Data was not split into further sub-elements at this stage due to multiples of the same types of element appearing for some METARs and SPECIs, for example if more than one layer of cloud was reported. These groups were split into a relational data structure with a separate linked table created for each field where multiple entries were possible. This is discussed in the section ‘Weather data extraction and validation into a relational data structure’ on page 20.

Some observation data existed in non-standard format, however, in many cases it was still possible to decode these records. To gather as much information as possible, a ‘leap-frogging style’ delimiting process was used to ensure that all information was retained, even if it was moved around within the text. An original reference was retained to ensure data integrity upon review.

The algorithm used a ‘regular expression’ based search algorithms to find an expected value, for example visibility or wind speed and direction. If data of this format is located within the text, this data is delimited with unique markers. If such data is not present, blank delimiting markers are inserted. These markers framed the start and end of each field of interest, allowing this data to be easily selected in subsequent steps, and blank entries to be added if these were not identified.

Additionally, the algorithm allows for invalid data prior to an expected value, for example if a particular row has data that is out of sequence or contains a typographic error. This invalid data is moved to the end of the text row for later review, and a Null entry is placed in that column. The algorithm then moves to the next column of expected data for assessment.

For every year of data, the following processes were applied to transform the raw METAR, SPECI and TTF text into a single flat file (table) of fields.

---

[5] Rainfall was extracted from within the remarks section.

1. Import raw text file
2. Remove all tab fields (Hex 09)
3. Remove all blank rows
4. Check for invalid row formats (those rows not starting with YYYYMMDD HHMM)
    a. Move all invalid rows up to the end of the previous row (as these seemed to be used for comments) and store these for later review and assessment
    b. All correctly formatted rows formatted to start with a single integer: YYYYMMDDHHMM
5. Prepare ‘identifier’ columns of data
    a. Remove known data discrepancies
        i. TTF included but not required and other non-standard coding
        ii. Some zulu times in reports missing day of month
            - Correct using date and time data from start of row identified in step 4 above
6. Identify and delimit data columns with tab fields
    a. Delimit ‘identifier’ columns with tabs according to the following format:
        i. YYYYMMDDHHMM M|S (METAR|SPECI) (AWS) ICAOCode DDHHMMZ?
    b. Identifier column names respectively are:
        i. DateTime, ReportTypeCode, ReportType, ICAOCode, ZuluTime
    c. Move data from invalid rows from the end of the text into a discrete column between DateTime and ReportTypeCode
7. Post processing of formed columns
    a. Zulu time has a Z added to the end if it was not present in the text
8. Extract and delimit data from column ‘report’ into grouped elements
    a. ‘Report’ column remains in an unaltered form for reference and validation of any changes, with an additional duplicate column created to be manipulated as required
    b. Known ‘error’ flags inserted by a BoM analysis were removed
    c. Additional white spaces were removed
    d. Develop regular expressions and extract and delimit data into specific grouped elements. Note:
        i. See Appendix A, Table 34 on page 95 and Table 35 on page 95 for the regular expressions and code used
        ii. If a data element is not found, a Null field is inserted and delimited in its place
        iii. A data search from left to right is conducted and updated progressively
        iv. Attempts are made to remove erroneous data for later review by moving this to the ‘error’ column created in 6(c) above
        v. Multiple groups of RVR, weather phenomena, cloud groups, recent weather and windshear are accounted for in this extraction
9. Export formed data as a tab delimited file with the ICAO code and dates in a standardised file name to allow autonomous retrieval

Sample of raw text field prior to extraction

```
20090210 1800 M METARAWS YBBN 101800 21002/05KT 9999 6ST016 6SC047 23.3/20.7
1005.3 RMK RF00.0/000.0/000  <<<  TTF:INTER 1800/2100 3000 RA BKN010 BKN020 OVC100

Error text
```

Sample of text after extraction

200902101800 *[TAB]* Error text *[TAB]* M *[TAB]* METARAWS *[TAB]* YBBN *[TAB]* 101800Z *[TAB]* 21002/05KT 9999 6ST016 6SC047 23.3/20.7 1005.3 RMK RF00.0/000.0/000 <<< TTF:INTER 1800/2100 3000 RA BKN010 BKN020 OVC100 *[TAB]* 210 *[TAB]* 02 *[TAB]* 05 *[TAB]* 9999 *[TAB]* *[TAB]* *[TAB]* *[TAB]* 6ST016 6SC047 *[TAB]* 23.3 *[TAB]* 20.7 *[TAB]* 1005.3 *[TAB]* *[TAB]* *[TAB]* *[TAB]* *[TAB]* *[TAB]* RMK RF00.0/000.0/000 *[TAB]* TTF:INTER 1800/2100 3000 RA BKN010 BKN020 OVC100

Delimitation of TAF data into a tab delimited flat file

TAF data was extracted using a similar approach to that of METAR, SPECI and TTF data above. However, only the unique identifiers for each row were delimited, specifically: date and time, report type, ICAO code and UTC time. The body of TAF reports was unaltered with the exception of separating the remarks section.

At the completion of the delimiting process for all records, TAF and TTF records were delimited with the same level of information extracted, allowing very similar processes to be used to extract data from both types of forecasts as described in ‘Weather data extraction and validation into a relational data structure’ on page 20. Table 2 shows the groups delimited by the process described below.

Table 2: TAF data flat file column names

Field type	Field name
Unique identifier group	Date and time
Non-standard text	Notes and error text
Unique identifier groups	Report type
	ICAO Code
	Issue Time
Forecast groups	Original forecast
	Converted ICAO forecast
	TAF remarks

In the same way as METAR, SPECI and TTF records, some TAF data existed in a non-standard format. The process used to delimit TAF data into a flat file, including small corrections to the formatting are described as follows:

1. Import raw text file
2. Remove all tab fields (Hex 09)
3. Remove all blank rows
4. Check for invalid row formats (those rows not starting with YYYYMMDD HHMMSS)
    a. Move all invalid rows up to the end of the previous row (as these seemed to be used for comments) and store these for later review and assessment
    b. All correctly formatted rows formatted to start with a single integer: YYYYMMDDHHMMSS
5. Identify and delimit data columns with tab fields
    a. Delimit ‘identifier’ columns with tabs, identifier column names respectively are:
        i. DateTime, ReportType, ICAOCode and ZuluTime
    b. Move data from invalid rows from the end of the text into a discrete column between DateTime and ReportType
6. Post processing of formed columns
    a. Zulu time has a Z added to the end if it was not present in the text
7. Extract and delimit data from column ‘report’ into grouped elements:

        a. ‘Report’ column remains in an unaltered form for reference and validation of any changes, with an additional duplicate column created to be manipulated as required
        b. Known ‘error’ flags inserted by a BoM analysis were removed
        c. Additional white spaces were also removed
        d. Validity periods, written with a space HHMM HHMM replaced with HHMM/HHMM as per ICAO format
        e. Develop regular expressions and extract and delimit data into specific grouped elements. Note:
            i. See Appendix A, Table 34 on page 95 and Table 35 on page 95 for the regular expressions and code used
            ii. If a data element is not found, a Null field is inserted and delimited in its place
            iii. A data search from left to right is conducted and updated progressively
            iv. Attempts are made to remove erroneous data for later review by moving this to the ‘error’ column created in 5(c) above
8. Export formed data as a tab delimited file with the ICAO code and dates in a standardised file name to allow autonomous retrieval.

Importation and validation of delimited weather data

A stored procedure was created to import the delimited weather data files (described above) into an SQL database.

These were imported directly from the delimited flat files created using a Microsoft PowerShell script, as noted in the section ‘Formation of weather flat data files’ on page 9. The standard file names of these flat files (as noted in step 9 for METARs, SPECIs and TTFs and step 8 for TAFs also in section ‘Formation of weather flat data files’) were selected by nominating start and end years, the ICAO code for the airports, and the root directory where the flat data files were stored. Data columns were the same as those created for the weather flat files.

To enhance confidence in data integrity, quality control and to prevent files being imported more than once, a log entry was created in a separate table for each file that was imported. A logical check was performed for every nominated file to be imported to ensure that the same file could not be loaded more than once.

However, to allow flexibility, the algorithm was designed to allow individual locations to be regenerated without affecting data from other locations. This was done in anticipation that data from some locations may contain errors or gaps requiring re-loading, while preserving any analyses performed at all other locations. This also allowed all data for analyses to be stored in one database, which was considered to make data tracking more efficient, analyses between locations easier, and provide a higher level of data integrity.

Gross error checks

To provide confidence in the imported flat data files mentioned above, an assessment algorithm was developed. This was designed to check two different aspects of the fields used as identifiers for both the TAF, METAR, SPECI and TTF flat files, as follows:

1. Check that recorded flat file names containing ICAO codes and data name matches all reports and forecasts within the flat file
    a. This was performed by extracting elements of the file name stored in the import log file mentioned above and ensuring that this information matched every row in the imported table
2. Correct column formats
    a. Date and time field length
    b. Report type field matches one of eight known types of forecast or report

c. ICAO code field contains only the Australian Y designator followed by 3 letters
d. Issue time field is of correct length

The algorithm prints specific feedback of the results of these error checks, either confirming that no errors were identified or identifying errors and the specific nature of the error in each case. If an error is identified, a table is produced with the erroneous data to enable troubleshooting.

Time period gap analysis

A stored procedure was developed to help with assessing if there were any significant periods without data that may influence error and uncertainty in the data analysis. The time between the release of subsequent reports and forecasts was evaluated for selected start and end years. Time gaps greater than 1 hour for METARs, SPECIs and TTFs, and 7 hours for TAFs were displayed for evaluation for every location imported.

Both the gross error checks and time period analysis can be performed at any point of data extraction.

Formation of instrument approach procedure flat data file

Aircraft instrument approach data was entered into a single Microsoft Excel data sheet from each pdf document as noted in the ‘Data sources’ section on page 1.

To prepare the data for extraction into a relational data structure, airport approach data was organised to store one row for every instrument approach at the selected locations. The column headings of this data file are shown in Appendix A, Table 37 on page 101. Figure 27 on page 103 contains the data used in both example columns in Table 37, and shows where some of this information is drawn from.

The first eight columns in Table 37 from the runway name up to the long approach name contain details specific to the instrument approach plate allowing unique identification. The ICAO code and runway name were used as the basis for linking airport approach data to a runway database and weather data (as discussed in the section ‘Evaluation of weather reports and forecasts’ on page 48).

There were four lookup tables explicitly linked to the approach data structure.

- Approach type
- ILS category
- Multi-variant design (MVD) category
- Aircraft performance category

The MVD and aircraft performance category tables are discussed in ‘Compilation of aircraft specification flat data file’ on page 16. The approach type and ILS category tables are discussed below.

A new row was created in the flat data file for every possible combination of approach type, ILS category and MVD category. Aircraft performance category data was inserted into the flat table structure under applicable columns suffixed with one of these four categories. This can be seen in the ‘Example airport approach data entry’ on page 15.

Approach type and ILS category

The approach type table (shown in Table 47 ‘Appendix B Lookup tables created for weather analysis’ on page 106) contains a list of ten categorised approaches that were identified at Australian airports. Each row in this table also corresponds to the 10 rows at the end of the aircraft specifications table (shown in Table 36 on page 101), and discussed in the section ‘Compilation of aircraft specification flat data file’ on page 16.

Seven ILS categories were defined[6] (shown in Table 46 page 105). These were used to link to applicable instrument approach procedures and aircraft capabilities.

The values shown in Table 46 are nominal values as defined by ICAO, however, the published instrument approach procedure decision altitude and visibility minima are used when assessing landing minima criteria for each aerodrome.

The ILS category was not explicitly published in some instrument approach procedures, mainly for category I instrument approach procedures (CAT I procedures). In these circumstances, the ILS category was determined using the following algorithm, based on the ICAO Manual of Aerodrome Standards[6].

- If DH = 0 and RVR < 50 Set Category III c
- Else If DH < 50 and 50 <= RVR < 200 Set Category III b
- Else If DH < 100 and RVR < 350 SET Category III a
- Else If 100 <= DH < 200 and RVR < 550 SET Category II
- Else If DH >= 200 or Visibility >= 800 or RVR >= 550 SET Category I

For ILS CAT I procedures, other aerodrome and aircraft factors may work to make these limits more stringent, such as if HIAL is not available as written in AIP 1.5 4.7.3.a, b and c. However, only weather factors were taken into account at the time of writing.

Example airport approach data entry

The approach plate shown in Figure 27 shows a matrix of landing and alternate minima requirements in the lower part of the figure. This shows an 'RNAV GNSS (RNP)' approach (recorded in Table 37 as approach type Id '3' as referenced in Table 47) that has two levels of required navigational performance (0.2 and 0.3 nautical miles) published. These are treated as two separate approaches, and this example focuses on an RNP of 0.2.

The ILS category does not apply in this case and is recorded with an 'Id' of '7', 'Not applicable' as shown in Table 46. However, in cases where more than one ILS category applies, a new row was created for each category.

Figure 27 also shows different criteria presented for three multi-variant design (MVD) groupings, MVDN, MVD2 and MVD4. For the purpose of this example, MVDN and MVD2 are shown in example 1 and example 2 respectively of Table 37. These are recorded with MVD category Ids of 2 and 3 respectively as per Table 3 below.

Once these rows were created based on all possible combinations of the approach type, ILS category and MVD category, the flat data file was populated with landing and alternate minima visibility and ceiling data.

For example, the RNP (0.2) MVD2 approach (shown in Figure 27 and Table 37 example 2 column) applies to two aircraft performance categories – C and D. The decision altitude is reported as 356 feet with a minima visibility of 900 metres for both categories. These were recorded under the minimum descent altitude and minimum visibility columns suffixed with 'C' and 'D'. All applicable data fields, such as alternate minima criteria and circling minima were also populated with data. Any data fields not applicable remained unfilled, for example Figure 27 shows that this approach is not applicable for category 'A' and 'B' aircraft and as such no data fields suffixed with an 'A' or 'B' were populated for this approach.

---

[6] Information from definitions for 'Instrument Runways in Section 1.1 of ICAO Manual of Aerodrome Standards for Aerodrome Design and Operations based on Annex 14, Third Edition, July 1999

Compilation of aircraft specification flat data file

Aircraft specifications were compiled into a single flat data file for use in determining the instrument approaches that a selected aircraft could fly, and subsequently the most precise and least precise approaches possible, given reported weather conditions. This was performed by matching aircraft specification data and runway approach data criteria, and is discussed in depth in section ‘Instrument approach procedure extraction into a relational data structure’ on page 46.

Aircraft parameters were documented and tabulated for each selected aircraft where possible. There were three main purposes for determining aircraft performance data, as follows:

- to define the aircraft limitations with regard to crosswind and tailwind landings
- to derive two aircraft performance categories for the purpose of matching with published instrument approach criteria
- to define aircraft performance category
- to define multi variant design category
- to develop a framework to allow the assessment between aircraft of different characteristics and capabilities.

The focus of analysis was on the reliability of aerodrome weather forecasting. To assess this, it was assumed that the Boeing 737-800 aircraft was capable of utilising all aerodrome instrument approach aids up to ILS category III b, and GLS. This made the aerodrome, rather than the aircraft the limiting factor in approach selection to assess the best possible approach that an aerodrome could offer. The determination of the aircraft approach, and multi-variant design categories are discussed below. ILS categories are discussed further in ‘Approach type and ILS category’ on page 14.

The flat data file documenting the list of parameters used to determine these categories are shown in Appendix A, Table 36 on page 101. Three additional lookup tables showing the multi-variant design (MVD), aircraft performance and ILS category values are shown in Table 3 on page 18, and Appendix B Lookup tables created for weather analysis, Table 45 and Table 46 on page 105 respectively.

In addition to aircraft specifications, the aircraft specifications table also contained binary fields nominating each type of instrument approach the aircraft model would be expected to be equipped to fly. For example, if the aircraft was expected to be able to fly an RNAV GNSS approach, this would be recorded as 1.

In some cases, particular aircraft specification values could not be determined. To approximate a missing value, the median for all other aircraft in that category were used in lieu. This approximation was not expected to significantly affect the analysis for most parameters, due to a landing criteria being based on coarse aircraft categories rather than specific flight dynamics. For reference, Table 36 on page 101 also shows data used for the Boeing 737-800 and the median values used when Boeing 737-800 data was not available.

Determining the aircraft performance category

According to definitions in the AIP[7], the performance category is defined in bands of airspeeds on approach to land. The airspeed used was the expected speed when crossing the landing runway threshold, referred to as V_{at}. This was calculated by the following equation:

$$V_{at} = \max\left\{Landing\ configuration\ stall\ speed\ (V_{SO})\times 1.3 \middle| Stall\ speed\ at\ 1\ g\ clean\ (V_{S1g})\times 1.23\right\}$$

If neither of these stall speeds were found, the reference speed on landing (V_{ref}) was used instead. This was calculated at the maximum landing weight where possible to represent a worst-case

---

[7] AIP Australia 17 NOV 2011 ENR 1.5 paragraph 1.2.1

scenario with higher airspeeds. Once these values were determined, the aircraft performance category was determined from the published speed ranges, as defined in the AIP[7], and shown in Table 45 on page 105.

Due to median values being calculated within each aircraft performance category to represent missing values, median values were not used in the determination of the aircraft performance category.

Determining the aircraft multi-variant design (MVD) category

The multi-variant design (MVD) category was used when calculating particular approaches requiring a specified navigational performance (RNP). This was based on a number of factors to categorise aircraft into one of four different MVD groups:

- MVDR – Regional airliners
- MVDN – Narrow body airliners
- MVD2 – 2 engined wide body airliners
- MVD4 – 4 engined wide body airliners

A number of aircraft parameters were collected to determine the MVD category for each aircraft. Limits were set for each category in ascending order based on general understandings for each category, which are shown below in Table 3.

The upper limit of the regional airliner group was based on the maximum aircraft range being less than 1,500 nautical miles. This was determined using the range of aircraft expected to be in this category, being pre-dominantly used in these types of ‘short-haul’ operations, however the definition is somewhat arbitrary in this analysis. Reported ranges from less than 1000 nautical miles to above 1700 nautical miles are recognised for these operations[8][9].

The next major threshold used was defining whether the aircraft had a wide-body. This was determined in accordance with the definition used by the Cambridge Aerospace Dictionary as having more than one aisle, necessitating being at least 4.72 metres wide. In addition, it was assumed that a maximum of six seats would be used.

As already noted, this table was used as a ‘lookup table’ in the primary aircraft table. Aircraft were determined to be in a category by assessing each category in ascending order, for example if an aircraft has a maximum range above 1,500 nautical miles, has 1 aisle, has less than 7 seats abreast and is under 5 metres wide, it will be categorised as a narrow body (MVDN) airliner. The Boeing 737 – 800 model used in the study at the time of writing has been categorised in the MVDN group.

---

[8] The Cambridge Aerospace Dictionary 2004 defines short-range transport as not exceeding 1,200 nautical miles in cruising conditions, short-haul transport as 1,000 statute miles (869 nautical miles)

[9] Her Majesty’s Revenue and Customs of the United Kingdom Excise Notice 550 quoted 2,000 statute miles (1,738 nautical miles) as the upper limit of short haul type flights.

Table 3: Aircraft multi-variant design (MVD) categories

Multi Variant Design Id	MVD Category	MVD Category Name	Restrictions applied				
			Maximum Width (metres)	Maximum Seats Abreast	Maximum Aisles	Maximum Range (nm)	Maximum Engines
1	R	Regional	5	6	1	1500	N/A
2	N	Narrow body	5	6	1	N/A	N/A
3	2	Wide body 2 engine	N/A	N/A	N/A	N/A	2
4	4	Wide body 4 engine	N/A	N/A	N/A	N/A	4

Relational data structure formation, extraction and validation

A relational data structure was considered the most desirable method for storing weather forecast and report data. The considered benefits of using a relational database structure include:

- allowing efficient storage of entities with multiple values within a group
- operations can be performed on a very large number of records
- data can be combined together for analysis based on conditions
- data integrity is maintained
- availability of the software.

The database developed for the analysis was set up to allow every item evaluated to be traced back to the original source documents, providing transparency and enabling data quality checks to be performed with relative ease.

One of the disadvantages of using an SQL relational database was considered to be related to numerical functions and calculations, where complex calculations across rows of data can be time-consuming, some requiring customised functions to be developed. In addition, conducting in-depth analysis and presentation of results was required to be output into another program, adding time to the analysis. Microsoft Excel was used for this purpose.

An empty relational database was created using SQL to contain all data required for analysis. The database framework was formed prior to any data being extracted, and was progressively populated with extracted data and the output of simulations and analyses.

An SQL stored procedure was created to develop the data framework. This procedure was the ‘child’ of the stored procedure which was used to manage the creation of the data framework and also import weather data used for the analysis, as described in ‘Importation and validation of delimited weather data’ on page 13. This data framework is stored in a single nominated database, and is created if it does not exist, or is nominated to be reset.

In comparison to a flat data file table structure, using a relational database allows one field to be nominated to hold all of that type of data, rather than requiring multiple, identical field types. It was considered that this would enable a significantly less cumbersome analysis with fewer coding errors, in particular for calculations such as determining aggregate figures, and querying across different rows, rather than columns.

The following two sections describe the data framework formation for processes used to extract each individual source of data prior to being combined together.

Formation of relational weather data framework

The following section describes the formation of a weather data structure to store all details of METAR and SPECI reports, and TAF and TTF forecasts. The discussion is limited to extracting the weather data itself, prior to it being combined with any other data. Further discussion on combining data is contained within the section ‘Evaluation of weather reports and forecasts’ on page 48.

METAR and SPECI report data structure

The data structure for METAR and SPECI reports consisted of one parent data table (METAR decode), seven child tables and two lookup tables. These tables are listed below, and are also described in the following section.

- Metar decode
    - Cloud cover
    - Ceiling estimate type (lookup)
    - Visibility estimate type (lookup)
    - Runway visual range
    - Phenomena
        - Condensed phenomena
    - Recent weather
        - Condensed recent weather
    - Runway windshear

Each child table of the parent METAR decode table was formed for each field with more than one possible input, such as cloud cover.

TTF data structure

To achieve the maximum fidelity possible, the data structure was formed to allow TTFs (and TAFs discussed below) with multiple forecast segments, such as FM, INTER and TEMPO to be extracted into separate data rows.

Furthermore, where the multiple forecast elements exist, such as different layers of cloud, separate data tables were required to describe these, for example one table includes details of every individual cloud layer reported or forecast, which in turn contains a link to the forecast segment that it is contained in. This follows the same general process as described in METAR and SPECI report data extraction below, with additional tables being created to expand grouped data.

Three ‘child’ level tables, one lookup table and one ‘parent’ level table were created in conjunction with the TTF data extraction, as follows:

- Delimited data table
- TTF decode
- Cloud cover
- Weather phenomena
- Condensed weather phenomena
- Segment time group convention Id (Lookup)
- Turbulence

TAF data structure

The overall process followed similar process to that described for TTFs, with each forecast segment being split into separate rows and linked to the delimited data table.

There were five grouped TAF fields which were extracted into three ‘child’ level tables, and three ‘parent’ level data tables. In addition, two references of the same lookup table (for the overall TAF and TAF segment timing elements) were inserted into the TAF decode table. This structure was as follows:

- TAF decode
- Weather phenomena
- Condensed phenomena
- Clouds
- TAF time group convention and segment time group convention (lookup)
- Turbulence
- QNH
- Temperature

Weather data extraction and validation into a relational data structure

Following the delimitation of data into the formats listed above, the weather forecast and report flat files were imported and extracted into the SQL data framework using the stored procedures developed for this purpose as described above. A number of measures were implemented to validate the data to provide confidence in the nature of the source data.

The following section describes the processes applied to weather data up to the point, but prior to the evaluation and simulation processes where this data was combined with aircraft and runway approach data.

The two delimited tables imported above were required to be extracted into a relational data structure for comparison and analysis of individual elements to take place. Data from the METAR, SPECI and TTF delimited data table was split into two of three main weather data tables, one for meteorological reports (METARs and SPECIs), and one for TTF data. The third data table was extracted from the delimited TAF data table.

The following sub-sections describe the processes used to extract, derive or estimate the weather information used in this analysis.

Figure 2: Overview of procedures used to extract delimited weather data into a relational data structure

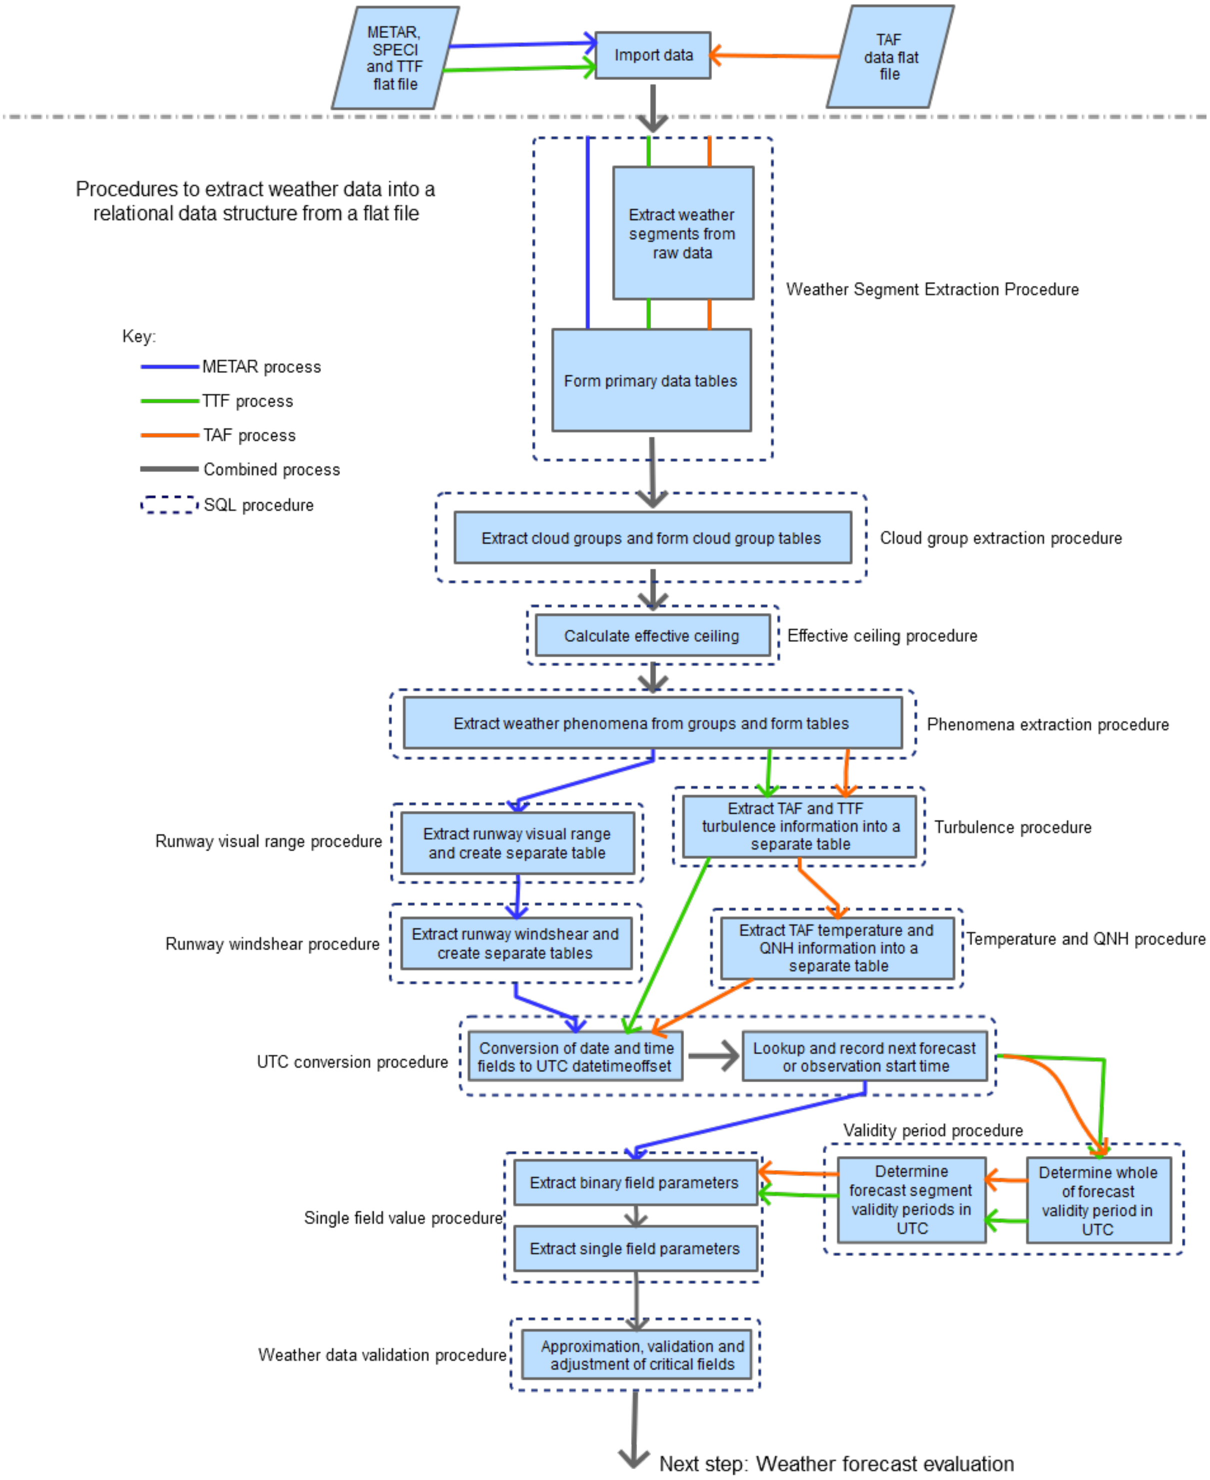

Formation of primary data tables – (Extraction of TAF and TTF weather segments, and METAR and SPECI data rows)

The following section describes the formation of the primary (parent) data tables within the TAF, TTF and METAR/SPECI data structures,as described in 'Formation of relational weather data framework' on page 19. These tables have links to the original source data flat files and are referenced in each of their child tables.

METAR and SPECI report data extraction

METAR and SPECI report data was extracted and stored in a data structure to enable comparison with each element of a forecast.

The vast majority of this data was drawn from the METAR, SPECI and TTF delimited data table, with the exception of two fields referencing runway data described in ‘Instrument approach procedure extraction into a relational data structure’ on page 46. Table 4 shows the processes applied to generate data in the main data table (METAR decode) from the visibility and cloud group fields of the delimited data file, as described in ‘Delimitation of raw METAR, SPECI and TTF data into a delimited flat file’ on page 9. A complete list of all mapped fields is contained in Appendix A, Figure 22 on page 96.

The heading ‘METAR, SPECI and TTF delimited data table’ on the left of Table 4 shows the column names of the delimited data file. In addition, the left side table also contains examples of typical, but not all of the types of data that may be contained in these fields. The column labelled ‘Data map number’ is used as a reference to show where data is sourced from in the METAR decode table.

The right side of Table 4 lists the fields (column names) of the primary table used to compare forecasts to METAR and SPECI reports throughout the analysis. The data contained in the output text column shows the resultant data from the sample text column displayed in the left hand side.

The column labelled ‘Data from’ shows the data sources used for each field sourced from the delimited data table. These correspond to the ‘Data map number’ for each field in the (left hand side) delimited data table. The column labelled ‘Calculation type’ lists codes indicating the level of processing applied to populate each field with data. Table 5 contains descriptions for these codes.

Table 4: Sample of data processes between delimited weather report data and extracted METAR table

METAR, SPECI and TTF delimited data table			METAR decode table				
Data map number	**Column name**	**Sample text**	**Data from**	**Calculation type**	**Column name**	**Output text**	**Additional table(s)?**
13	Visibility	2000	13, 17	A	Visibility	2000	
			13	E	CAVOK	0	
			13, 17	A	VisibilityEstimateId	1	Yes
14	MinVisibility	1000NE	14	E	MinimumVisibility	1000	
			14	E	VisMinDirection	NE	
			14	D	VisMinDirection_degrees	45	
17	CloudGroups	VV014; 2ST006 7SC038 1CB032	17	E	VerticalVisibility	1400; NULL	
			13, 17	A	CeilingEffective	1400; 3800	
			13, 17	A	CeilingOctas	NULL; 8	
			17	U	CloudGroups	VV014; 2ST006 7SC038 1CB032	Yes
			17	E	SKC	0	
			17	E	NSC	0	
			17	E	NCD	0	
			13, 17	A	CeilingEstimateId	1	Yes

There were 23 columns where no change was made to the data from the delimited data file. These were included to allow a single reference table of extracted data. Five of these unaltered columns were grouped data stored to be split into additional data tables, as indicated in the ‘Additional tables’ column. The creation of additional data tables from the grouped data is explained in ‘Extraction of grouped data into additional tables’ on page 25.

The five derived data columns consisted of:

- a plain text description of both recent and observed weather phenomena
- the minimum visibility direction in degrees
- a conversion of the time field stored as an integer to a date time offset field to allow a more robust analysis of time.

Not all unaltered columns were extracted in the same process. This was to allow efficiency when building all tables, starting with only the grouped columns, which required further extraction, and finishing with the single value parameters.

Of the eight columns containing extracted information, five related to binary condition flags (CAVOK, SKC, NSC, NCD and SkyObsc). If an indicator such as 'CAVOK' was contained in the delimited file, a '1' was inserted in the corresponding field, whereas the absence of this text would be stored as '0'. The remaining extracted information related to minimum and vertical visibility. These were mainly related to and stored in either visibility or cloud information in the delimited data file as shown in Table 4.

Table 5: Calculation types from delimited to extracted data

Calculation type	Code	Description
Un-modified direct copy	U	No change to data from delimited file
Derived from data	D	Information from data field(s) used to determine a value directly with minimal processing. For example the explicit storage of time in UTC rather than an integer, or storing a compass point recorded as NE as 45 degrees.
Algorithm based	A	A more complex determination is used to determine the value, possibly involving multiple data sources. This process may involve approximations using the information available.
Extracted directly from field	E	Data fields are split into sub-components for storage in separate fields. For example minimum visibility stored as 1000NE is divided into separate fields: 1000 metres visibility, and NE direction.

Five fields were produced (calculation type labelled 'A' in Table 4) using two algorithms relating to the calculation of visibility and ceiling for use in the analysis. The determination of these fields often required interpreting information contained in flags or different types of information. In some cases, insufficient information was available to derive these fields, and approximations were made. This process is described in 'Determining and validating weather ceiling and visibility' on page 41.

TTF data extraction

TTF data was extracted from the same delimited data table as for METARs and SPECIs as described in the section immediately above. As noted, unlike METARs and SPECIs, TTF data did not have a '1 to 1' relationship with the delimited data file, as multiple rows were produced to split each TTF into 1 row per forecast segment, requiring additional processes to be developed.

The extraction of grouped data and time group data is described in the sections starting on pages 25 and 33 respectively.

Table 6 below shows an example of 1 row of data from the delimited data table being divided into 3 separate TTF forecast segments. Critical column names in the extraction are also shown. This process is generally described in Figure 3 on page 25.

The first row (labelled row 1) of each decode consists of an extraction of the reported conditions up to the time that the next segment became active. For example, in Table 6 this 'segment' is only active for 1 minute due to the second segment validity period starting from 1950, with the TTF being released at 1949. In cases where 'NOSIG' is reported (indicating that there is no significant change from the observation expected), only one segment is produced from the TTF release time for 3 hours in duration.

Each segment contains precise times that are determined using the time code immediately after the segment indicator, and the forecast release time (indicated by the year date time field). For example, the third segment listed commences with ‘FM2230’, indicating the forecast segment from 2230 UTC. When combined with the information from the ‘year date time’ field, it can be determined that this corresponds to the same time on 6 June 2009. An in-depth discussion for the determination of forecast timing fields is contained in the section ‘Derivation of TAF, TTF, METAR and SPECI release and segment change times into universal coordinated time (UTC)’ on page 33.

Table 6: Extraction of TTF elements into separate rows from delimited table

METAR, SPECI and TTF delimited data table (Parent table)		**TTF decode (Child table)**			
Column name	**Row 1**	**Column name**	**Row 1**	**Row 2**	**Row 3**
		TTF Decode Id	99948	99949	99950
BoM Bulk Metar Data Id	91486	BoM Bulk Metar Data ID	91486	91486	91486
Year Date Time	200901061949	Year Date Time	200901061949	200901061949	200901061949
TTF Group	TTF:FM1950 20010KT 9999 DZ BKN015 FM2230 20012KT 9999 BKN020	TTF Segment Type	TTF METAR	FM	FM
		TTF Segment Order	1	2	3
		TTF Segment Text	TTF:	1950 20010KT 9999 DZ BKN015	2230 20012KT 9999 BKN020
		Segment Start Time UTC	2009-01-06 19:49:00 +00:00	2009-01-06 19:50:00 +00:00	2009-01-06 22:30:00 +00:00

A data mapping table showing all column names in the TTF table extracted from the parent delimited data table is shown in Appendix A, Figure 23 and Figure 24 on pages 97 and 98. Figure 2 on page 21 shows the order in which each of the elements were extracted for the TTF data tables.

Turbulence and remarks data removed prior to extraction

To minimise variability when separating the TTF groups into segments, turbulence and remarks information was extracted from the TTF group text prior to the separation algorithm. This was done for several reasons:

- It was identified that a number of rows in the TTF decode table were extracted incorrectly as a result of turbulence text being very similar to the standard from ‘FM’ forecast segment being contained in the main body text, rather than the remarks section of text.
- The remarks section may apply to the all TTF segments, necessitating this text being linked to all segments.
- The turbulence section has potential timing elements which may be independent of the forecast segments, meaning that this required a link to be developed directly to the delimited data table (discussed in Extraction of forecast turbulence from TAFs and TTFs on page 31).

The TTF report prior to the modification was preserved with turbulence and remarks data in the ‘OriginalTTFtext’ field as listed in Figure 24 on page 98. Turbulence and TTF remarks data was also preserved in fields in the TTF decode table.

TAF data extraction

TAF data was drawn from a specific delimited data file, as described in ‘Delimitation of TAF data into a tab delimited flat file’ on page 12.

The extraction of TAF data into multiple elements (as shown by Table 6 above) followed the same process as for TTFs with the exception that turbulence and remarks groups were not required to be extracted. This was because it was identified that all turbulence data was contained in the remarks groups, which were already extracted in the delimitation process.

The extraction of all grouped TAF data, and the determination of timing elements is discussed in depth in ‘Extraction of grouped data into additional tables’ and ‘Derivation of TAF, TTF, METAR and SPECI release and segment change times into universal coordinated time (UTC)’ on pages 25 and 33 as per METAR, SPECI and TTF extractions.

A data mapping table showing all column names in the TAF table extracted from the parent delimited data table is shown in Appendix A, Figure 25 and Figure 26 on pages 99 and 99 respectively. Figure 2 on page 21 shows the order in which each of the elements were extracted for the TAF data tables.

Extraction of grouped data into additional tables

For the purpose of referential integrity, grouped data was sourced from the already formed METAR decode table. This allowed the unique identifier columns of the METAR decode column to be inserted into each of the child tables as shown in Figure 3 below, establishing a parent – child relationship.

Figure 3: Creating a child data table from a grouped text field

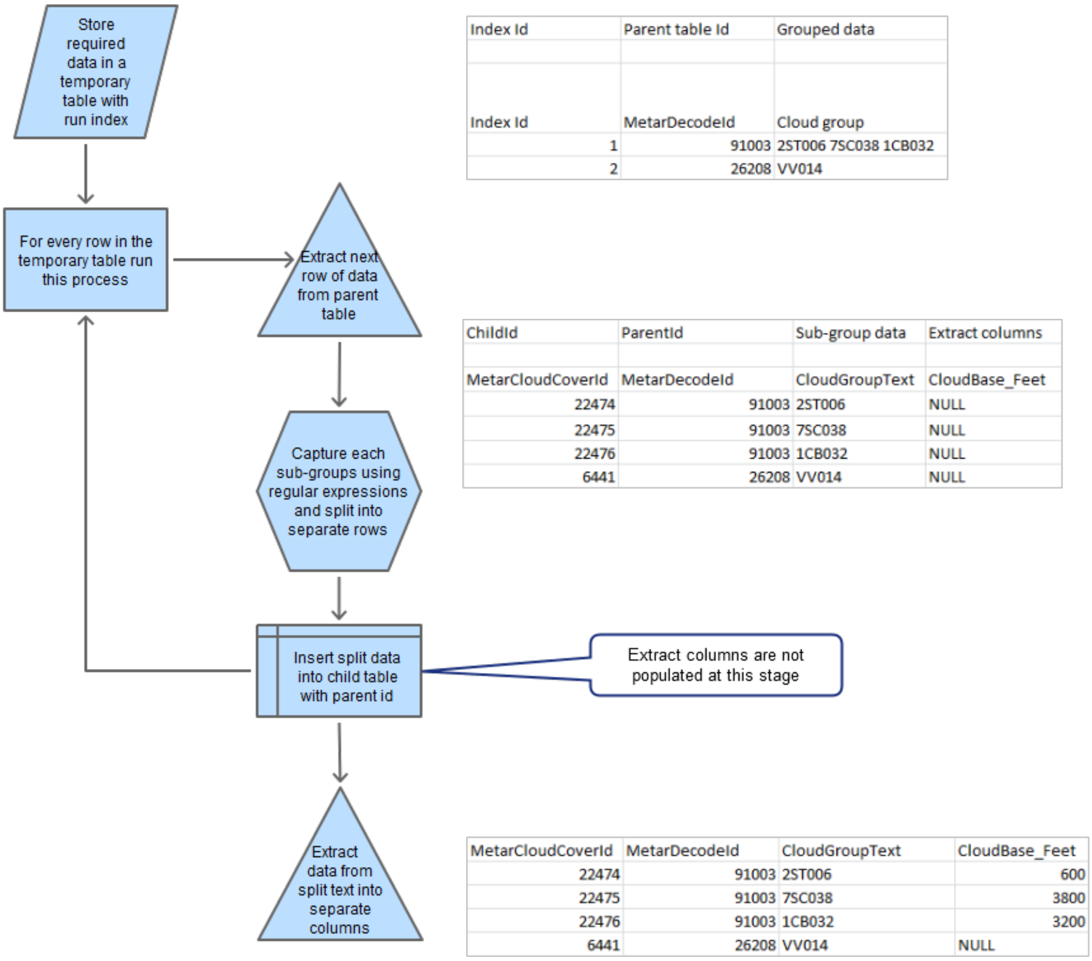

Index Id	Parent table Id	Grouped data
Index Id	MetarDecodeId	Cloud group
1	91003	2ST006 7SC038 1CB032
2	26208	VV014

ChildId	ParentId	Sub-group data	Extract columns
MetarCloudCoverId	MetarDecodeId	CloudGroupText	CloudBase_Feet
22474	91003	2ST006	NULL
22475	91003	7SC038	NULL
22476	91003	1CB032	NULL
6441	26208	VV014	NULL

MetarCloudCoverId	MetarDecodeId	CloudGroupText	CloudBase_Feet
22474	91003	2ST006	600
22475	91003	7SC038	3800
22476	91003	1CB032	3200
6441	26208	VV014	NULL

Extraction of cloud groups

Cloud groups were the first of the grouped data extracted from the parent METAR decode table. Using the process shown in Figure 3, cloud group data was first split into the component groups with each group then derived separately.

Two lookup tables were created to describe the cloud elements in further detail, specifically, the cloud coverage codes, and the reported classification of cloud type. These tables are shown in Table 38 and Table 39 of Appendix B Lookup tables created for weather analysis on page 101.

A sample of the cloud group extraction from a single row in the parent METAR decode table is shown in Figure 4 below. The two columns on the left show the data used by the parent table, the unique identifier (METARDecodeId), and the grouped data to be extracted (Cloud Group). The four columns on the right show the resultant data from the left side columns that were inserted into all of the columns as listed off the MetarCloudCover table with the exception of MetarCloudCoverId. This is a sequential unique primary key that is automatically generated for every row created in MetarCloudCover.

Figure 4: Sample extraction of cloud group text into METAR cloud cover table

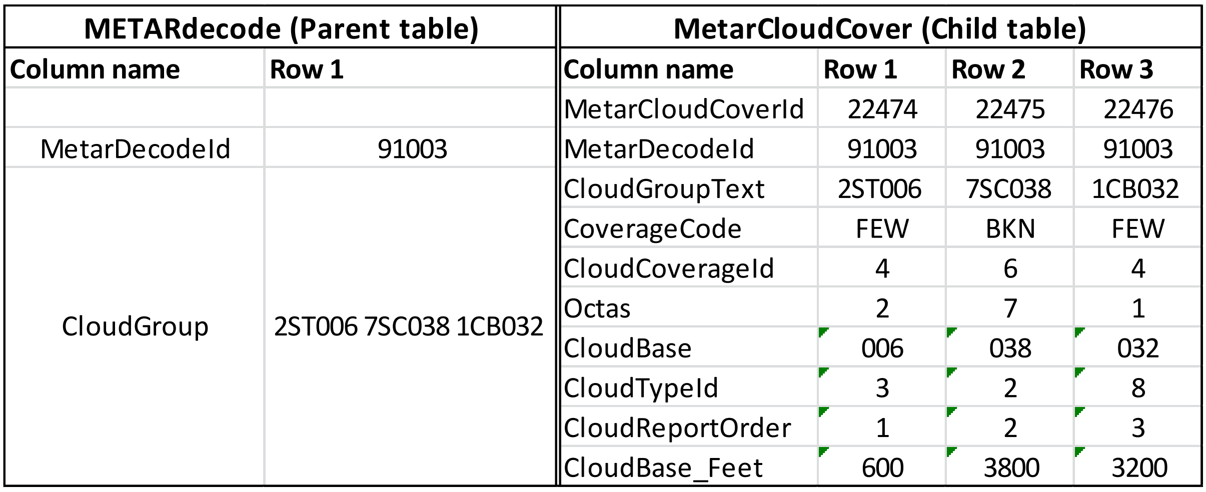

METARdecode (Parent table)

Column name	Row 1
MetarDecodeId	91003
CloudGroup	2ST006 7SC038 1CB032

MetarCloudCover (Child table)

Column name	Row 1	Row 2	Row 3
MetarCloudCoverId	22474	22475	22476
MetarDecodeId	91003	91003	91003
CloudGroupText	2ST006	7SC038	1CB032
CoverageCode	FEW	BKN	FEW
CloudCoverageId	4	6	4
Octas	2	7	1
CloudBase	006	038	032
CloudTypeId	3	2	8
CloudReportOrder	1	2	3
CloudBase_Feet	600	3800	3200

As per the rules stipulated in the AIP, worst-case cloud coverage scenarios were applied. For example, if FEW clouds were reported, this would be derived as two oktas to allow summation of coverage. This is in line with the AIP rules stipulating FEW + FEW = SCT, FEW + SCT = BKN, etc. This is discussed further in the Calculation of the weather ceiling on page 43.

The algorithm accounted for circumstances where cloud coverage was reported in oktas, or as a coverage code such as SKC or BKN, populating any missing data possible. For example, a cloud group with BKN038 would be recorded as seven oktas at 3800 feet, however the cloud type is not known.

Coverage codes reported as SKC (Sky clear), NCD (Nil cloud detected) and NSC (Nil significant cloud) are recorded as zero oktas for that row of data where identified. However, according to the definitions of the AIP, these values do not necessarily mean that there was no cloud. This is discussed further in ‘Calculation of the weather ceiling’ on page 43.

These extractions and derivations occurred in the following order:

1. Extract oktas where reported
    a. Derive cloud coverage codes from reported oktas
2. Extract coverage codes where oktas not reported
    a. Nominate oktas from cloud coverage codes from reported coverage codes

3. Extract reported cloud base
4. Extract cloud type where reported
    a. Lookup cloud type in lookup table (as shown in Appendix B Lookup tables created for weather analysis Table 38)
5. Lookup and insert cloud coverage id from cloud coverage codes formed in steps 1 or 2 above (lookup table shown in Appendix B Lookup tables created for weather analysis Table 39)
6. Convert reported cloud base to feet (multiply 3 digit cluster by 100)
    a. A small number of cloud base fields were reported in a 2-digit cluster, instead of the standard 3-digit cluster. In these cases, the 2-digit cluster was multiplied by 1000. As this produced a conservative figure (a higher cloud base) and due to the small number of 2 digit clusters, it was considered that this would not adversely affect the analysis.

This data table formed the primary method in determining the observed ceiling as described on page 43.

Extraction of weather phenomena (including recent weather)

Weather phenomena and recent weather were extracted using the same stored procedure, as shown in Figure 2 on page 21. This procedure was used to populate four of the seven child tables of the parent METAR decode table, consisting of:

- weather phenomena (primary table with split data for each phenomena sub-group)
- recent weather phenomena (primary table with split data for each recent weather phenomena sub-group)
- condensed weather phenomena (all weather phenomena in one row for each weather report or forecast)
- condensed recent weather phenomena (all recent weather phenomena in one row for each weather report or forecast)

Two lookup tables, weather intensity and weather vicinity, were created to further document weather phenomena descriptions. These tables are shown in Table 40 and Table 41 of Appendix B Lookup tables created for weather analysis on page 104.

The descriptions below are shown for the weather phenomena tables, however, the process followed and column names were identical for recent weather tables.

An example of this extraction showing how the data is split into the weather phenomena table is shown in Table 7 below. Condensed phenomena data is shown in Table 8. All phenomena tables contain one column for each of the 20 phenomena and 8 descriptors, however, only phenomena reported in the example are shown.

The weather analysis does not rely on reported weather phenomena (with the exception of thunderstorms) as performance criteria for landing or alternate minima calculations. Therefore, if phenomena were not reported explicitly or otherwise, they were assumed absent, even if NSW (Nil Significant Weather) is not forecast.

To extract weather phenomena data, the same general process was followed as per the cloud groups, and other grouped data, as illustrated in Figure 3 above.

Table 7: Sample extraction of weather phenomena group text into METAR Phenomena table

METARdecode (Parent table)		METARPhenomena (Child table)		
Column name	**Row 1**	**Column name**	**Row 1**	**Row 2**
		Metar Phenomena Id	2968	2969
Metar Decode Id	68612	Metar Decode Id	68612	68612
Weather Phenomena Groups	+DS VCSH	Weather Phenomena Text	+DS	VCSH
		Phenomena Decode	heavy dust storm	showers in the vicinity
		Intensity Id	3	5
		Weather Vicinity Id	0	1
		SH	NULL	1
		DS	1	NULL

Separating different groups of phenomena was important to enable the discernment of different weather codes, such as heavy, moderate or light precipitation, or if particular phenomena were only observed in the vicinity, and not directly overhead.

Documenting exactly which phenomena were reported or forecast overhead or in the vicinity was important when excluding conditions that probably did not affect an aircraft making a safe landing. For example, thunderstorms are more significant when overhead, than within 8 to 16 kilometres of the aerodrome.

A separate function was developed to extract weather intensity and produce a plain text description of the phenomena cluster. One of the uses of this information is the assessment of heavy precipitation on the runway as described in ‘Wind vector calculations and comparisons’ on page 62. It was deemed that a separate function was required due to this algorithm being used multiple times in the program and being more complex, because of the weather intensity indicators having potentially different meanings depending on the weather phenomena in the group.

To create separate rows and divide weather phenomena data, a new row in the child phenomena table was created wherever a space was present in between valid phenomena, as shown in the ‘weather phenomena text’ column output in Table 7.

Once these rows were created, the table was progressively populated with information derived from this field as follows:

1. Using the function created (FN_WxPhenDecode on page Y) by converting phenomena codes
   a. Insert plain text description
   b. Insert intensity Id field (as shown in Appendix B Lookup tables created for weather analysis Table 40 on page 104)
2. Lookup and insert weather vicinity Id field (as shown in Appendix B Lookup tables created for weather analysis Table 41 on page 104)
3. For every weather phenomena and descriptor
   a. Assess if each are contained in the text in turn
   b. If present, insert a ‘1’ into the corresponding column labelled for that phenomena
      i. For example if SH is present, insert a ‘1’ is inserted under the ‘SH’ column of the table
   c. If weather phenomena are not present, the field remained unchanged, as a ‘Null value’

In addition to the primary weather phenomena table listed above, a table (labelled ‘condensed phenomena’) was produced to summarise weather phenomena into a single row of data for each weather report or forecast. This table was developed to prevent cluttering in the main data table to store each of the 28 weather phenomena and descriptors (prior to the relational data structure being developed). Computer code was developed to store an indicator under each phenomena (such as ‘VC’ or ‘+’) if these were applicable, however at the time of writing this was not required to be utilised. The phenomena decode column contains a plain text description of the entire phenomena group for the weather report or forecast.

A sample of the condensed phenomena table drawn from grouped data in the METAR decode table is shown in Table 8 below. For comparison, this data is the same used in Table 7 above.

Table 8: Sample formation of condensed phenomena table

METARdecode (Parent table)		**METARCondensedPhenomena (Child table)**	
Column name	**Row 1**	**Column name**	**Row 1**
		Metar Condensed Phenomena ID	2922
Metar Decode Id	68612	Metar Decode Id	68612
Weather Phenomena Groups	+DS VCSH	Weather Phenomena Text	+DS VCSH
		Phenomena Decode	heavy dust storm with showers in the vicinity
		SH	1
		DS	1

A similar process was followed to create the condensed phenomena table to that performed for the primary phenomena table, with the exceptions that the weather intensity (step 1.a) and weather vicinity (step 1.b) fields could not be extracted.

Extraction of runway visual range from METAR and SPECI reports

Runway visual range data was reported for one or more runways at some airports, for example Sydney International and Melbourne International airports. Dividing this data for each reported runway was required to allow assessment against the landing minima criterion for particular instrument approaches. As for all data falling under the Extraction of grouped data into additional tables, the general approach depicted in Figure 3 on page 25 was followed.

Runway visual range data was split into one row for each runway reported. Individual elements were derived or extracted into a separate table. The column names and a sample of the extraction from grouped text are shown in Table 9 below. The algorithm for extraction was contained within a single stored procedure without the requirement for additional functions due to this process being used once per ‘run’ of the overall extraction algorithm.

The column labelled ‘Runway Id’ was created to establish a link between this data and a runway data table. This table is discussed in ‘Runway details’ on page 5, and allows runway approach data to be merged with reported weather conditions specific to each runway if available.

Table 9: Sample formation of runway visual range table from METAR and SPECI reports

METARdecode (Parent table)		METARRwyVisRange (Child table)		
Column name	**Row 1**	**Column name**	**Row 1**	**Row 2**
		Metar Rwy Vis Range ID	8	9
Metar Decode Id	115470	Metar Decode Id	115470	115470
RVRGroups	R16L/1200VP2000D R16R/0325N	Runway ID	1059	1061
		Rwy Vis Range Text	R16L/1200VP2000D	R16R/0325N
		Runway Number	16	16
		Parallel Runway	L	R
		RVR Greater Than Max	1	0
		RVR Less Than Mnm	0	0
		Sig Variation	1	0
		Ten Min Avg	NULL	325
		Mnm One Min Avg	1200	NULL
		Max One Min Avg	2000	NULL
		Distinct Tendency	Downward	Nil

The process for extraction and derivation of the data was as follows:

1. Directly extract runway details and flags
    a. Runway number
    b. Parallel runway code
    c. RVR quoted greater than maximum range of equipment (P)
    d. RVR quoted less than minimum range of equipment (M)
    e. Significant variation in RVR (V)
2. Populate runway visual range values
    i. If no significant variation is reported, populate ten minute average field with quoted value
    ii. If significant variation is reported, populate minimum and maximum one minute average values
3. Derive reported distinct value, downward, nil or upward
4. Lookup the runway Id from the runway data table (discussed on page 5) using the ICAO code and extracted runway details.

Extraction of runway windshear from METAR and SPECI reports

Runway windshear was extracted where these elements were identified in weather reports. This involved a similar approach taken to the ‘Extraction of runway visual range from METAR and SPECI reports’ on page 29, in that data was extracted into a table with one row per reported runway, and a link was made with the runway data table (discussed on page 5).

Runway windshear had not been used for analysis at the time of writing. However, it was considered that this data might be useful at some future stage. The process followed was steps 1.a, 1.b and 4 of the ‘Extraction of runway visual range from METAR and SPECI reports’ on page 29. A sample extraction for runway windshear, showing the column names for this table is shown in Table 10 below.

Table 10: Sample formation of runway wind shear table from METAR and SPECI reports

METARdecode (Parent table)		METARRwyWindshear (Child table)	
Column name	**Row 1**	**Column name**	**Row 1**
		MetarRwyWindshearId	1
Metar Decode ID	99206	MetarDecodeId	99206
Windshear groups	WS RWY23	Runway ID	840
		RwyWSText	WS RWY23
		RwyName	23
		RwyNum	23
		ParrallelRwy	NA

Extraction of forecast turbulence from TAFs and TTFs

Turbulence was extracted from the formed ‘turbulence groups’ columns in the primary TAF and TTF tables. A temporary data table was created with one row per TAF or TTF containing turbulence information. These two tables contained a link to the delimited data tables and the turbulence group text. The delimited table link was required instead of the parent TAF or TTF decode tables due to the timing of forecast turbulence being synchronised independent of TAF and TTF forecast segment timing.

Turbulence information was not used for the analysis at the time of writing, however, it was considered important that the data structure accommodated this for potential analysis. The extraction of turbulence groups followed the same general process as for all other grouped data, as illustrated in Figure 3 on page 25.

Table 11 shows the column names of the TAF turbulence table, and the links to the table directly created from the delimited data file, ‘BoMBulkTAFData’, instead of the TAF decode table as noted above. The TTF turbulence table is formed in the same way, with the exception that TTFs are linked to the METAR data table, as described in ‘TTF data extraction’ on page 23.

The maximum height and severity codes were extracted once each turbulence segment had been split into the separate rows of the table (also shown in Table 11). The worst-case turbulence severity was recorded if turbulence of varying severity is forecast. At the time of writing, the start and end time algorithms for turbulence were not completed, which may be required for future analysis of this data.

Table 11: Sample formation of turbulence table from a TAF

BoM Bulk TAF Data (Parent table)		TAF Turbulence (Child table)		
Column name	**Row 1**	**Column name**	**Row 1**	**Row 2**
		Taf Turbulence Id	134	135
BoMBulkTafDataId	2499	BoMBulkTafDataId	2499	2499
Turbulence Groups	FM131800 MOD/SEV TURB BLW 5000FT TILL140300 FM140400 MOD TURB BLW 5000FT	Turbulence Text	FM131800 MOD/SEV TURB BLW 5000FT TILL140300	FM140400 MOD TURB BLW 5000FT
		Severity Code	SEV	MOD
		Max Height_FT	5000	5000
		Start Time UTC	NULL	NULL
		End Time UTC	NULL	NULL

Extraction of forecast temperature and QNH values from TAFs

Forecasted temperature and QNH (sea level pressure) were extracted from TAFs into two separate tables, ‘TAFTemperature’ and ‘TAFQNH’. In a similar way to turbulence, this information was not critical to analysis at the time of writing, although was extracted to complete the framework for further studies.

Temperature and QNH data was joined with the delimited data file ‘BoMBulkTAFData’ (following the same process as turbulence data) rather than the ‘TAF decode’ table. This was due to the forecast time of each of the four forecast temperature and QNH values in each TAF being independent of the TAF segment start times.

Table 12 and Table 13 show a sample extraction and column names for each of these tables. Each temperature or QNH value forecast is split into a separate row. In addition to this, the orders that the values appear in the TAF are recorded to allow the calculation of the times when required. At the time of writing, the start and end times of these values was not calculated, as it was not required for the initial analysis undertaken.

Table 12: Sample formation of temperature table from a TAF

BoM Bulk TAF Data (Parent table)		TAF temperature (Child table)				
Column name	**Row 1**	**Column Name**	**Row 1**	**Row 2**	**Row 3**	**Row 4**
		TafTemperatureId	1	2	3	4
BoMBulkTafData Id	1	BoMBulkTafDataId	1	1	1	1
Temperature Groups	T 15 17 22 26	Temperature	15	17	22	26
		Display Order	1	2	3	4
		StartTimeUTC	NULL	NULL	NULL	NULL
		EndTimeUTC	NULL	NULL	NULL	NULL

Table 13: Sample formation of QNH table from a TAF

BoM Bulk TAF Data (Parent table)		TAF QNH (Child table)				
Column name	**Row 1**	**Column Name**	**Row 1**	**Row 2**	**Row 3**	**Row 4**
		TafQNHId	1	2	3	4
BoMBulkTafData Id	1	BoMBulkTafDataId	1	1	1	1
QNH Groups	Q 1007 1010 1010 1009	QNH	1007	1010	1010	1009
		Display Order	1	2	3	4
		StartTimeUTC	NULL	NULL	NULL	NULL

		EndTimeUTC	NULL	NULL	NULL	NULL

Derivation of TAF, TTF, METAR and SPECI release and segment change times into universal coordinated time (UTC)

Determining the UTC for each overall forecast validity period was considered important to allow the precise synchronisation of data between weather reports and forecasts, and other peripheral data sources such as sunrise and sunset times. As this was a formalised time code in the computer program, it also allowed the duration of time to be measured between points with relative ease.

Varying levels of processing were required to derive complete times for these weather forecasts and reports. Weather reports (METARs and SPECIs) required minimal calculations, whereas TAFs and TTFs required a more complex algorithm to be developed. An overview of these processes is shown in Figure 5 below.

Prior to extraction of forecast validity periods, an exact reference time in the highest fidelity available was required. This was due to forecast segment validity periods being abbreviated, and there being insufficient information to determine the exact time without reference to other parts of the TAF or TTF. The reference time used was the forecast and observation release date time field, as noted in paragraph 4.b for METARs, SPECIs and TTFs on page 11 and paragraph 4.b for TAFs on page 12 in the ‘Formation of weather flat data files’ from page 9.

There were four fields requiring a complete time group in the TTF decode table, corresponding to the start and end times of the overall forecast and each individual TTF segment. TAFs also had these fields, with an additional fifth field corresponding to the end of the becoming (BECMG) segment if applicable.

Figure 5 shows the sequence of operations performed to populate all of the time fields in the primary TAF, TTF and METAR/SPECI tables. Each main process is labelled as a single digit on the left side of the diagram, with sub-processes labelled immediately to the right.

The red and green boxes labelled ‘Sub-routine: Time group convention’ and ‘Sub-routine: Time group extract’ indicate where additional algorithms were used to determine the values of the sub-process to which these are attached. These sub-routines are discussed in depth in ‘Identifying forecast time groups’ on page 35, and ‘Determining UTC time for TAF and TTF partial time elements’ on page 37.

Flags at the right side of the diagram show values that were used as function inputs, or values directly used to populate particular values. For example, process 2.1 on the determination of the TTF start time always uses the ‘release time’ as the ‘start time’. Flags attached to the function label ‘Time group extract’ indicate the reference time used in implementation of this function.

TTF validity

Validity of TTFs was assumed to be 3 hours after release. For that reason, the start time of all TTFs was assumed to be the release time (as described in ‘Converting the release time of forecasts and reports to UTC’ on page 34), and the end time of all TTFs was assumed to be 3 hours after this time. This information was calculated and recorded in the start and end time fields of all TTFs accordingly.

Figure 5: Overview of processes used to extract weather forecast and report times

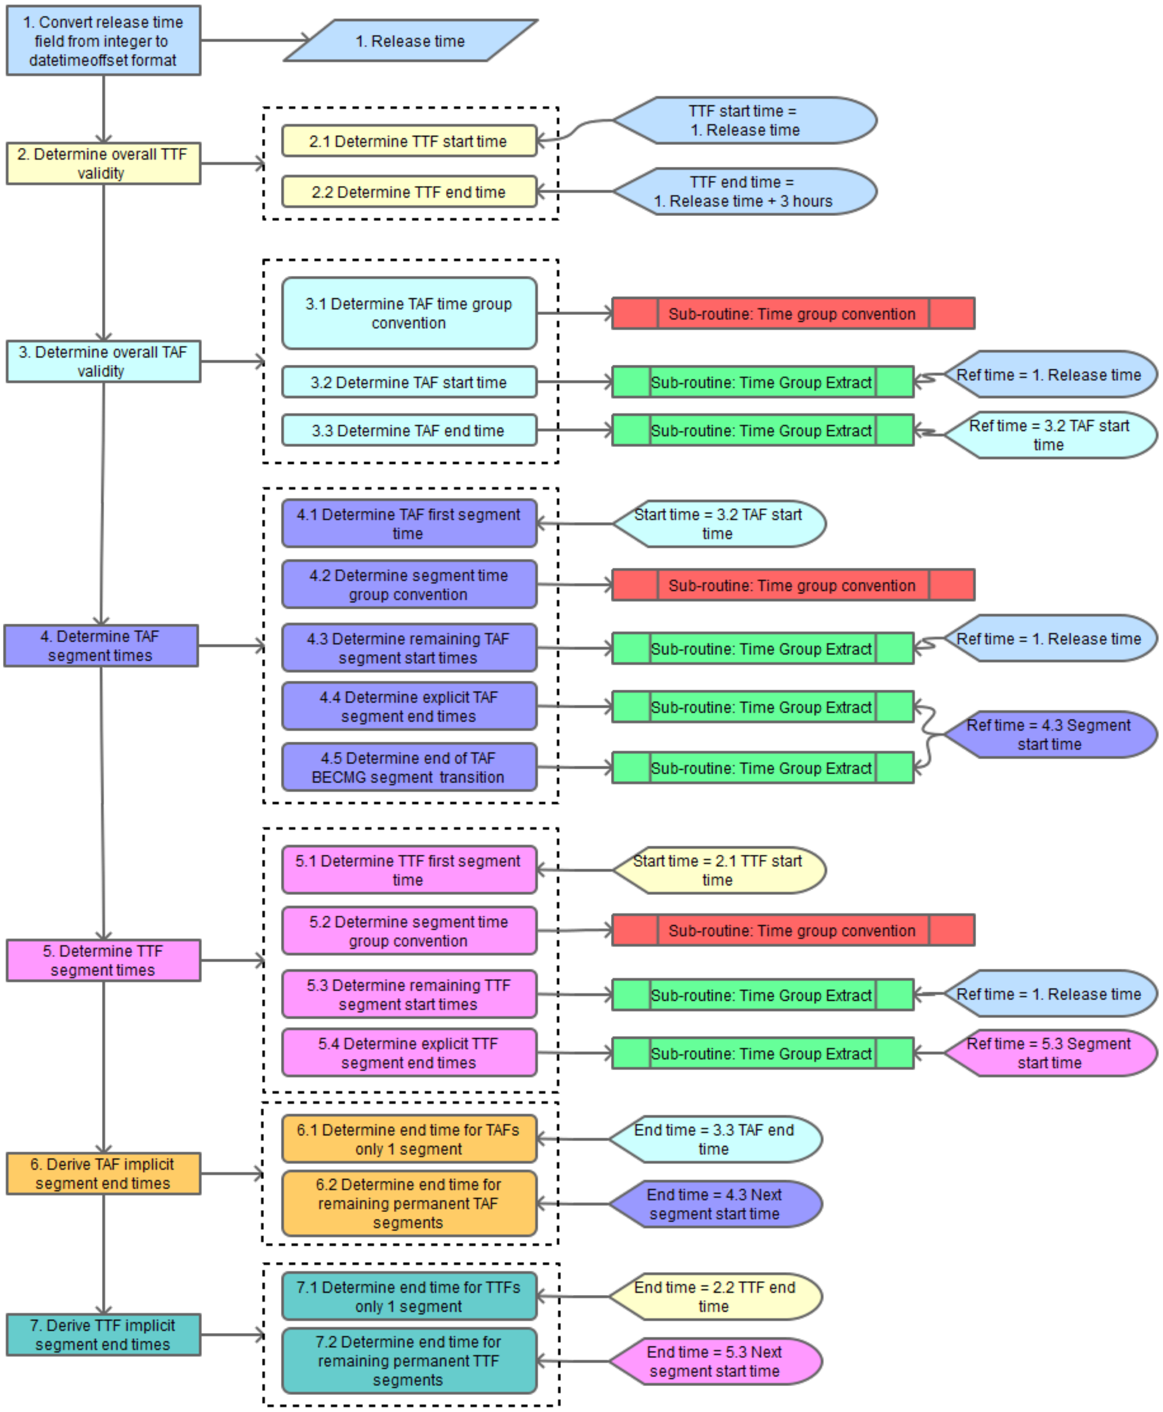

The following sections discuss the conversion and derivation of weather report times which were subsequently used as the basis for combining most of the data in the analyses.

Converting the release time of forecasts and reports to UTC

An algorithm was written to populate the ‘release date time UTC’ field of each of the TAF, TTF and METAR decode tables. For precision and to allow comparison with data with different reported time zone offsets, a ‘datetimeoffset’ field was chosen to store all time fields. Using the extracted date time field in the format YYYYMMDDHHMM for METARs, SPECIs and TTFs and YYYYMMDDHHMMSS for

TAFs, data was converted and inserted into the relevant field of these tables. Examples of this conversion are shown in Figure 6 below, and Appendix A Figure 22 on page 96. Note that the release time data for a TTF was drawn from the METAR or SPECI to which it was attached.

Figure 6: Conversion of time from integer to datetimeoffset format for a TAF

YearDateTime	ReleaseDateTimeUTC
20090117223646	2009-01-17 22:36:46 +00:00

For every TAF, TTF and weather report, the date and time of the next record released was recorded in the ‘Next…datetimeoffset’ field. This was particularly important for METAR and SPECI weather reports as this was used as the end time for combining these reports with TAF and TTF weather forecasts for comparison. The comparison process is discussed in section ‘Comparison of overlapping TAF and TTF forecasts with METAR and SPECI reports’ on page 70.

Identifying forecast time groups

Prior to converting partial time groups into UTC time, the time group formats of each whole of TAF, and TAF and TTF segments were identified and stored for reference. It was considered that discrete pre-identification of time groups provided three main benefits when compared with performing this operation in a single cluster of conditional computer code.[10] This method was considered beneficial as it provided:

- a method of verification of the time-based data elements prior to conversion
- for the assumptions on the format of every time based value to be explicitly documented
- a mechanism to standardise and streamline the calculation of a complete time by using the same code for each time group identified.

Figure 7 below illustrates the processes used to determine the time groups published in a particular TAF or TTF. To allow this process to be used multiple times, this was created as a discrete computer function, and run in the procedure several times with different selected inputs.

This function was developed to account for different abbreviated time groups being present depending on the type of forecast, forecast segment or a historical period in time. For example, at the time of writing, the ICAO standard for reporting the validity period of TAFs was an 8-digit time code, split with a slash ‘/’ in the middle. Historically, other abbreviated time groups have been identified, such as the 6-figure time group DDHHHH.

The dark blue circles in Figure 7 correspond to conditional branches in the computer function. These are:

- the type of forecast, TAF or TTF
- if the overall TAF validity period is to be decoded or a TAF or TTF segment times are to be decoded.

The yellow diamonds were an additional assessment of the segment type, if a segment decode was selected.

The orange diamonds indicate where are a limited number of assessments were made at the end of each branch of the decision tree, starting with the most probable number format expected in each case.

Some orange diamonds have an optional length for the number of digits, varying by 1-digit. This indicated where the omitting of the last digit was allowed. This was due to a number of records being identified as missing this digit, and it was expected that this would produce the most probable match. The decoding of these values is discussed in ‘Determining UTC time for TAF and TTF partial time

---

[10] The previous iteration of this code was written in this way.

elements' on page 37.One advantage of identifying the most probable match was considered to be the ability to decode the start time completely, even if the end time was incomplete.

Figure 7: Determination of time groups in aviation weather forecasts

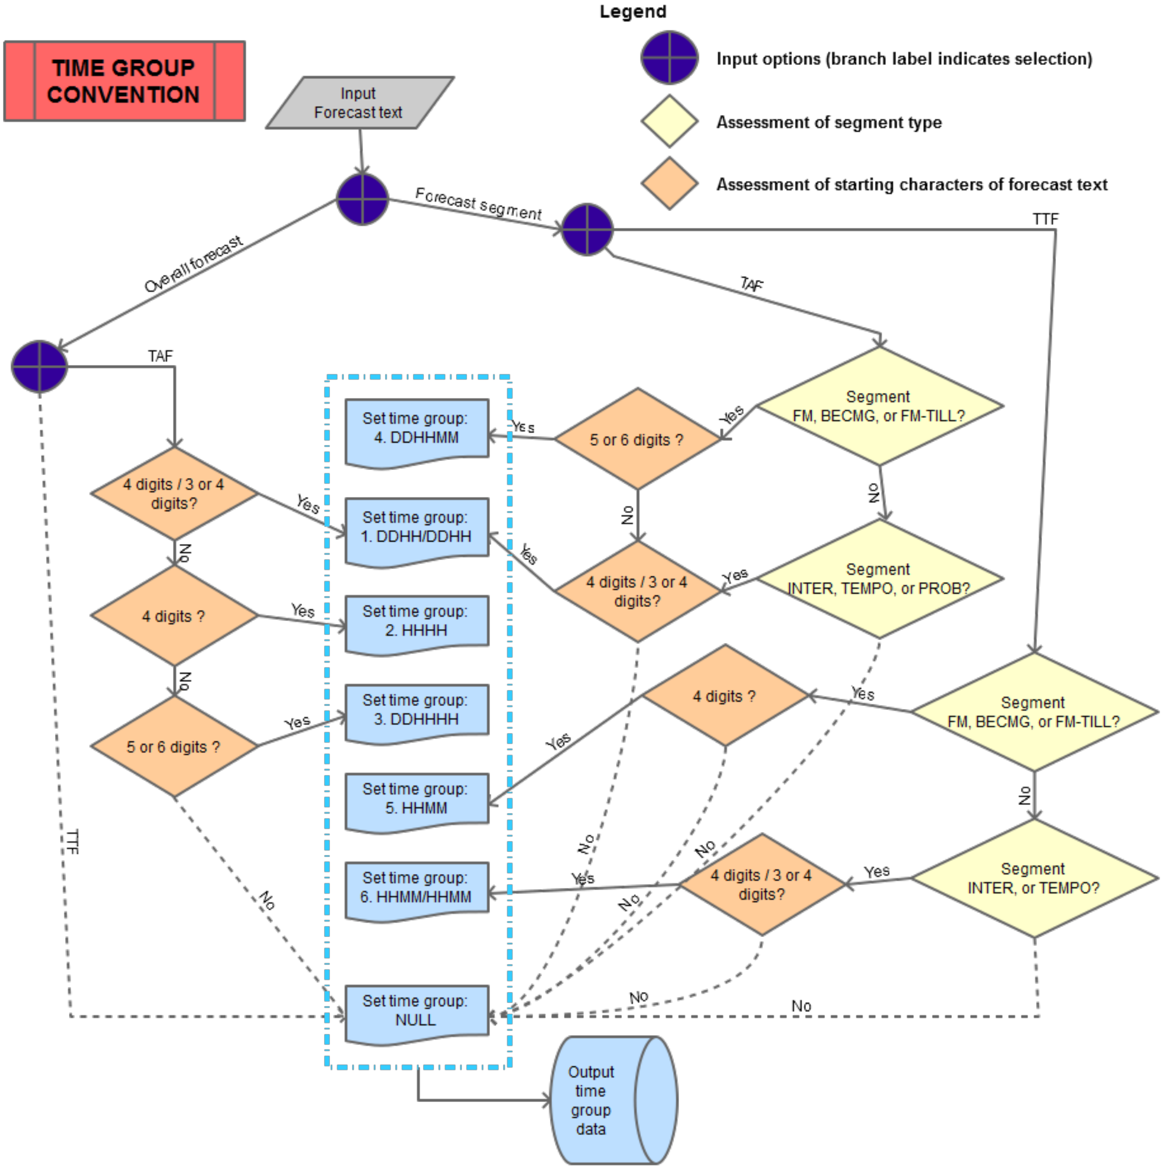

The function works by taking an input of text of either the entire forecast (for an assessment of the whole of forecast validity), or a forecast segment (for an assessment of that time group), and assessing the format of the first group of characters in the text field. A sample of this process is shown in Table 14.

Table 14: Sample output from function to determine forecast time groups

Input				Output
Forecast or Segment (F/S)	**Forecast type**	**Segment type**	**Text**	**Time group identified**
F	TAF	NULL	3118/0112 23010KT C…	DDHH/DDHH
S	TAF	FM	010400 21015KT C…	DDHHMM
S	TTF	TEMPO	2230/2400 6000 RA…	HHMM/HHMM

Appendix B Lookup tables created for weather analysis, Table 42 on page 105, shows the list of abbreviated time groups considered in the assessment. The numbered Id column in this table corresponds to the light blue output boxes in Figure 7, for example, the 8-digit time group ‘DDHH/DDHH’ corresponds to Id row 1 of Table 42.

Due to some time-groups in different types of forecasts and forecast segments having the same numerical format with different meanings (for example the 8 figure groups DDHH/DDHH and HHMM/HHMM as shown in Table 14 above), data was treated differently, depending on these elements. Specifically, the ‘time group convention Id’ columns of TTF decode table (Figure 24) and TAF decode table (Figure 26) were populated for records with values meeting a number format at the end of each decision branch. Records where these could not be identified remained ‘Null’, allowing immediate identification that the process was not successful. This also assisted in preventing the invalid conversion of a time group.

Start segments of TAFs and TTFs are not recorded, as these do not have explicit abbreviated time groups, and remain ‘Null’. Additionally, due to TTFs having no abbreviated time group for the overall validity period (due to the implicit validity period of 3 hours after release), the overall forecast time group for all TTFs also remains ‘Null’ as shown.

Determining UTC time for TAF and TTF partial time elements

Sub-routine overview

Using a complete time as a reference and the known time group format, complete UTC times were determined for the validity periods of overall TAFs and for forecast segments in TAFs and TTFs. Figure 8 below illustrates the main processes involved in the computer sub-routine to determine the partial time groups for all aviation weather forecasts in the analysis.

The sub-routine progressively ‘builds’ a time, one element at a time, using data from external inputs, and previous processes within the sub-routine.

The sub-routine uses the following inputs to achieve the conversion, as shown in the grey trapeziums in Figure 8:

- a complete reference time near the expected partial time
- whole of forecast or forecast segment text – depending on what time is being decoded
- time group information (described in ‘Identifying forecast time groups’ on page 35)
- a nomination of whether the start or end time is to be extracted.

This function was run five times when extracting times for every TAF and twice for extraction of partial time elements from every TTF, as shown in Figure 5 on page 34. The reference times used for forecasts are also shown in Figure 5.

Figure 8: Determination of complete UTC time from partial time fragment using a complete reference time in temporal proximity

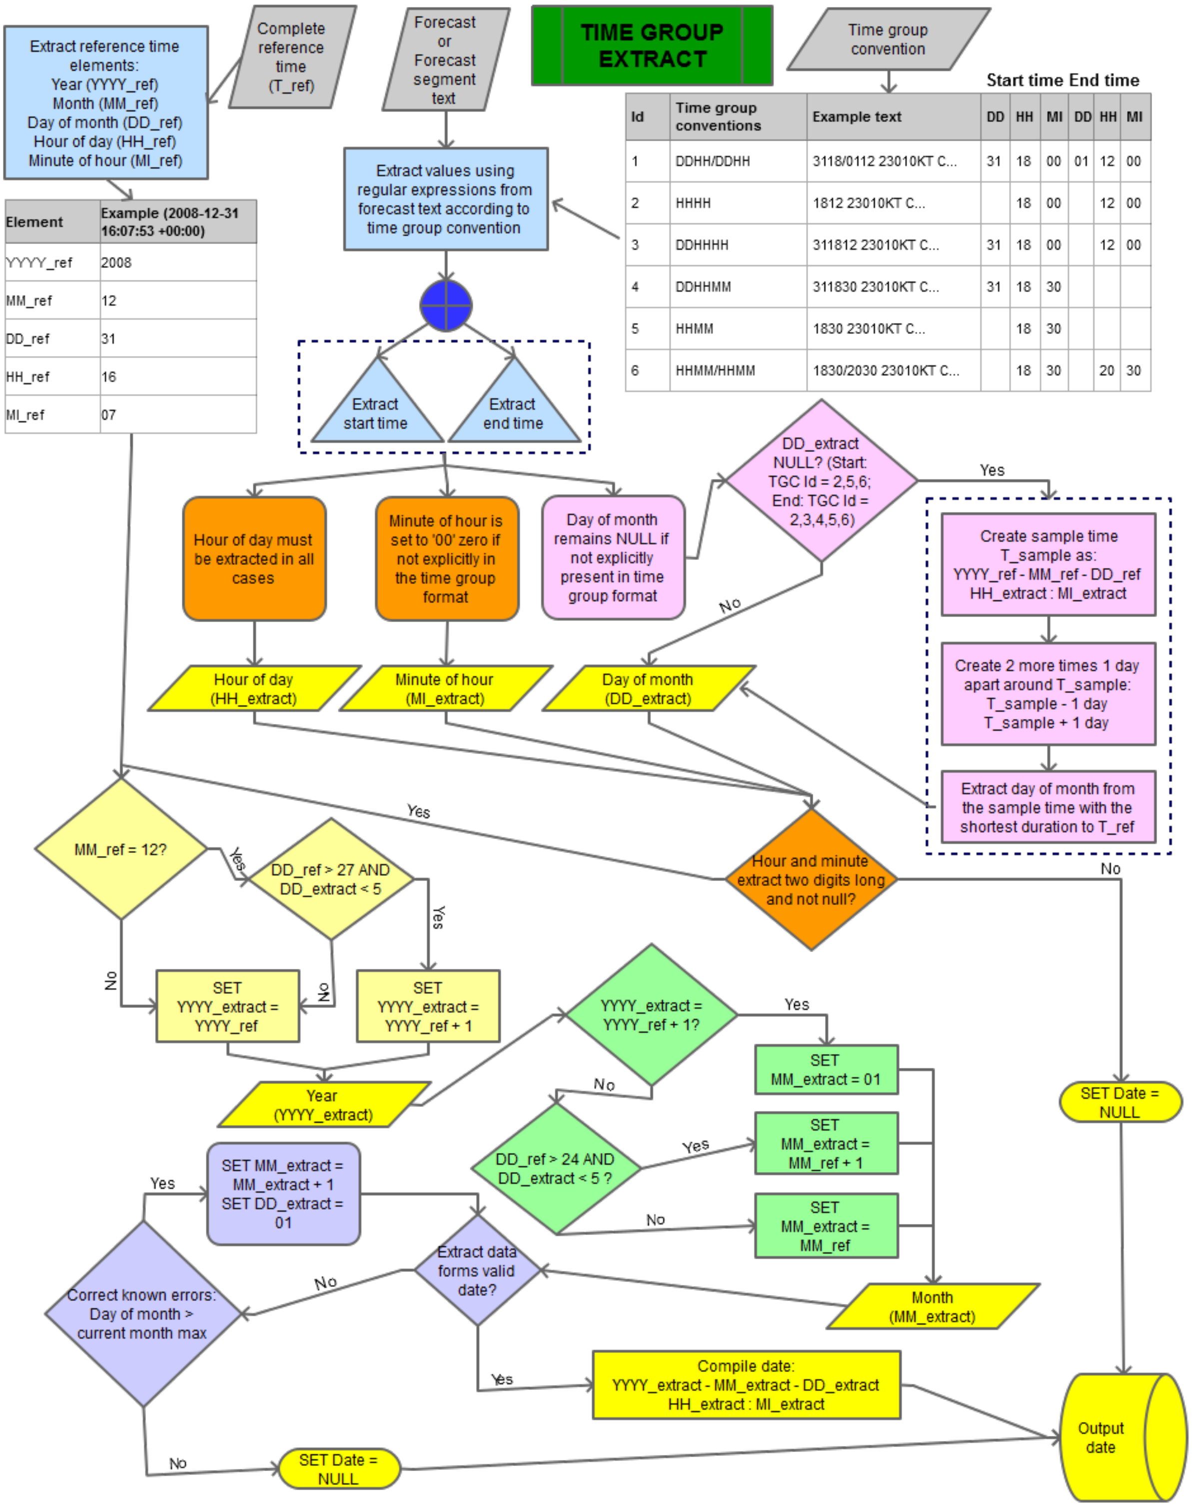

Element	Example (2008-12-31 16:07:53 +00:00)
YYYY_ref	2008
MM_ref	12
DD_ref	31
HH_ref	16
MI_ref	07

			Start time			End time		
Id	Time group conventions	Example text	DD	HH	MI	DD	HH	MI
1	DDHH/DDHH	3118/0112 23010KT C...	31	18	00	01	12	00
2	HHHH	1812 23010KT C...		18	00		12	00
3	DDHHHH	311812 23010KT C...	31	18	00		12	00
4	DDHHMM	311830 23010KT C...	31	18	30			
5	HHMM	1830 23010KT C...		18	30			
6	HHMM/HHMM	1830/2030 23010KT C...		18	30		20	30

The time group convention provided as a function input provides instructions for the extraction. An example of this extraction into day of month, hour of day and minute of hour is shown in the top right table of Figure 8.

Table 15 shows example final outputs of this sub-routine.

Table 15: Sample output from function to extract time from partial time fragments

Forecast type	Segment type	Function inputs				Output time
		Reference time	Forecast text	Time group convention	Start / End?	
TAF	TAF	2008-12-31 16:07:53 +00:00	3118/0112 23010KT CAVOK	DDHH/DDHH	Start	2008-12-31 18:00:00 +00:00
TAF	FM	2008-12-31 16:07:53 +00:00	010400 21015KT CAVOK	DDHHMM	Start	2009-01-01 04:00:00 +00:00
TTF	INTER	2009-01-28 13:00:00 +00:00	1430/1600 5000 TSRA SCT100…	HHMM/HHMM	Start	2009-01-28 14:30:00 +00:00
TTF	INTER	2009-01-28 14:30:00 +00:00	1430/1600 5000 TSRA SCT100...	HHMM/HHMM	End	2009-01-28 16:00:00 +00:00

Determination of day of month field when missing

In cases where the day of month was not explicitly in the time group convention, the reference time components, year and month were combined with the extracted hour and minute of day for use as a sample time. This is depicted in the pink boxes in Figure 8. Three sample times were produced 1 day apart as shown, and then compared to the reference time. The day of month for the time closest to the reference time was selected.

This process works for reference times up to twelve hours apart from the expected time, and allows for some cases where the forecast validity period is prior to the release time. It was expected that these ‘pre-dated’ forecasts arose from small delays in releasing a forecast due to provide coverage of a pre-determined period of time.

The allowance of only 12 hours either side of the reference time in determining the likely day of month could potentially result in an invalid time for the time group convention HHHH. This was attached to some older TAFs with validity periods longer than 12 hours, making the selection of the reference time more important. However, at the time of writing, none of these fields had been identified in the data analyses conducted, with only TTFs (with a maximum 3 hours validity) having omitted day of month fields. Therefore, improvements were not deemed a priority for the study at hand.

As shown in Figure 5 on page 34, reference times used for start times were the forecast release times, while the end times generally used the calculated start time. An improvement could be made as a future refinement of the algorithm, by progressively updating each segment using the previous segment as the reference time. Additionally, a forward-looking bias of possible reference times could be implemented to favour times after the reference time, while still allowing for forecast validity periods starting a nominated time prior to the release of the forecast.

Accounting for year and month transitions

Checks were made to assess if there was a change of month or year. This implementation is depicted in Figure 8 as green and light yellow shading respectively, using data from the provided reference time and the extracted day, hour and minute.

Error checking and data validation

Only valid date-time fields were output from the sub-routine (depicted in the purple shading in Figure 8). However, in cases where an invalid date was due to the forecast day being greater than the end of a month, this was corrected. For example, a day forecast as 29 February in a non-leap year was corrected to 1 March.

Additionally, extracted hour and minute data was required to be complete, containing 2-digits.

Extraction of single field parameters into primary METAR/SPECI, TAF or TTF tables

Binary field extraction

Descriptive flag fields were explicitly stored as a binary ‘1’ or ‘0’ column in each of the parent tables when these were identified or not present respectively. These fields were used as indicators of particular values in other fields when these were not recorded. For example, when the ‘CAVOK’ field was present in the absence of visibility and cloud information, the visibility was recorded as 9999 and the cloud base 4999. This is discussed further in section ‘Approximation of weather ceiling and visibility using weather condition flags’ on page 44.

Table 16 shows the descriptive flag fields that were stored as a binary field for each weather report or forecast. The use of these fields are discussed in ‘Determining and validating weather ceiling and visibility’ on page 41.

Table 16: Binary fields used in primary weather report and forecast tables

Binary field name	TAF	TTF	METAR/SPECI
Wind VRB	1	1	0
SKC	1	1	1
NSC	1	1	1
NCD	0	0	1
CAVOK	1	1	1
CLD:SKY MAY BE OBSC	0	0	1
NOSIG	0	1	0
NSW	1	1	0

Extraction of single field parameters

Weather parameters, such as wind speed and visibility where only one value per row in the primary table existed, were calculated after all other fields had been identified and extracted. This was required as some of these values were dependent on previous operations being complete, such as the identification of a binary field, or the compilation of the weather ceiling if data was available.

All of these values were unchanged from as they appeared in the delimited data tables, or directly derived, for example the simple conversion of the ‘S’ (South wind direction) to 180 degrees. These were extracted or derived using regular expressions into the primary rows of each parent table as follows:

- Metar decode parent table
- Wind direction, average wind speed, maximum wind speed
- Minimum visibility, Minimum visibility direction (compass points and degrees), Vertical visibility
- QNH, Temperature, Dewpoint
- Rainfall last 10 and 60 minutes, and since 9am
- TTF decode parent table
- Wind direction, average and maximum speed
- Wind shear group text
- TAF decode parent table
- Wind direction, average wind speed, maximum wind speed
- Wind shear group text.

Two values, visibility and ceiling, were determined using algorithms that were more complex. The assumptions and processes used for these two fields are discussed in depth in ‘Determining and validating weather ceiling and visibility’ below.

Determining and validating weather ceiling and visibility

Reported visibility and ceiling information was critical to the overall analysis of forecast reliability. This was due to these being the main elements used in published navigation thresholds to determine if the landing or alternate minima are breached at an aerodrome. This is described further in ‘Evaluation of weather reports and forecasts’ on page 48.

In cases where no explicit cloud or visibility information was contained in the report, a number of other information sources were used in lieu. In cases where these values could not be identified or approximated, these records were excluded and assumed safe conditions for the purposes of the analysis. For consistency in comparisons, the same algorithm was used for all weather forecasts and reports; for example, when no cloud information was explicitly reported and fields such as CAVOK are used instead. In some cases, expected data was not present, requiring a small number of records to be either excluded, or approximated if it was considered that sufficient data was available.

Table 17 below shows values inserted for visibility and ceiling for different published information in forecasts and reports. It also shows the order of preference that the computer algorithm follows when extracting this information, from explicit values of visibility of cloud data to averaging between different reports. Information of the highest level (lowest numbers presented first) in the hierarchy were used wherever possible.

Visibility or ceiling values were labelled N/A, not applicable, for weather elements not used in these respective calculations. The calculations of visibility and ceiling were performed independently (with the exception of slant visibility height calculations described on page 45), with the determination of one element performed, even if the other could not be calculated.

The flags NSW (Nil significant weather) and NOSIG (No significant change from weather report) were not used in these assumptions. This is because NSW was assumed to relate to the absence of weather phenomena and NOSIG was assumed to represent a continuation of reported conditions.

Table 17: Assumed visibility and ceiling in the absence of explicit values for elements in TAF, TTF, METAR and SPECI weather products

Order of preference	Value Category	Weather Element	Applicable weather products[11]	Nominal value		METAR Estimate Id	
				Visibility	Ceiling	Visi-bility	Ceiling
1	Explicit	Visibility or cloud data	METAR, TAF, TTF	As reported[12]	Reported cloud layers	1	1
2	Weather condition flags	SKC	METAR, TAF, TTF	N/A	9999	N/A	1
		NSC	METAR, TAF, TTF	N/A	4999	N/A	1
		NCD	METAR	N/A	4999	N/A	1
		CAVOK[13]	METAR, TAF, TTF	9999	4999[14]	1	1
3	Explicit	Vertical visibility	METAR	As reported	As reported	2	1
4	Ceilometer and visibility sensor anomaly	CLD:SKY MAY BE OBSC	METAR	As reported or 999	3 degree slant visibility height[15]	3	4
5	Explicit	Remarks section data	METAR	As reported	As reported	3	3
6	Average	Average between reports	METAR	Average of surrounding reports	Average of surrounding reports	4	5

The calculation of visibility and weather ceiling values as shown in Table 17 are described in the sections below. The columns under the label ‘METAR Estimate Id’ correspond to the columns VisibilityEstimateId and CeilingEstimateId in the METAR decode table. These were inserted for every value to indicate the amount and type of approximated data. Appendix B Lookup tables created for weather analysis Table 43 and Table 44 on page 105 show the descriptions for each of these approximations. This was the final data inserted into the weather data structure.

Calculation of weather ceiling and visibility using explicit values

The weather ceiling was determined for each METAR, SPECI, TAF and TTF using explicit values where possible from the respective cloud group tables (as described in ‘Extraction of cloud groups’ on page 26).

Visibility and minimum visibility fields were directly extracted from the published values, as shown in **Error! Reference source not found.** on page **Error! Bookmark not defined.**. Minimum visibility was not approximated or determined using any other means.

---

[11] Due to the first segment of a TTF being the METAR body text, METAR values are applied in these cases for the first period of the TTF, including where NOSIG is forecast.

[12] This includes explicitly recording minimum visibility in a separate field, as shown in Table 5 on page 20.

[13] Values based on CAVOK definition from AIP GEN 3.5 section 12.13 dated 29 MAY 2014

[14] Although the CAVOK definition requires that clouds are also above the lowest minimum sector altitude in a 25 nautical mile radius, the focus of these calculations are to align the values of forecasts and observations overhead the aerodrome in a consistent way.

[15] Discussed in ‘Ceilometer and visibility sensor data anomaly (METARs and SPECIs only)’ on page 37.

Calculation of the weather ceiling

The weather ceiling was determined for each METAR, SPECI, TAF and TTF using explicit values where possible from the respective cloud group tables (as described in ‘Extraction of cloud groups’ on page 26).

Values of calculated ceiling and the number of oktas were recorded in the parent METAR, TAF and TTF tables in this process.

The process involved the following steps for every weather forecast or report with explicit cloud information:

1. Arrange cloud data from lowest to highest cloud base
2. Create a cumulative sum of the number of oktas in the ordered cloud data
3. Assess the cumulative octa sums
    a. If cumulative sums exist above 4 oktas
        i. The cloud base in the first row above 4 oktas is recorded as the ceiling in the parent table
        ii. The number of oktas at this layer are also recorded in the parent table[16]
    b. If cumulative sums exist between 1 and 4 oktas
        i. The cloud base is recorded as 4999 feet
        ii. The number of oktas is recorded as calculated
    c. If the cumulative sum is 0[17]
        i. The ceiling is reported according to the weather condition flag as shown in Table 17 above (approximations discussed in ‘Approximation of weather ceiling and visibility using weather condition flags’ on page 44)
        ii. The number of oktas is reported as 0

Table 18 shows examples from conditions 3.a and 3.b above.

Table 18: Sample outputs from explicit calculation of weather ceiling

TAF cloud group table				Cumulative sum of Oktas	TAF decode table	
Taf Decode Id	**Cloud Group Text**	**Oktas**	**Cloud Base (Feet)**		**Effective Ceiling**	**Oktas**
574	FEW020	2	2000	2	4999	4
	FEW080CB	2	8000	4		
775	SCT020	4	2000	4	4000	8
	BKN040	7	4000	11		
	BKN100	7	10000	18		

Process 3.b.i above notes that cloud base is recorded as 4999 feet when the coverage was calculated between one and four oktas. This was done as a flag to indicate that the cloud base was at least above this level, and to be in line with the CAVOK definition (that clouds are at least above 5000 feet) [14]. This allows for observations of CAVOK conditions to be compared with a forecast where the visibility may be expected to be degraded, and no effective ceiling expected. This was considered conservative due to this weather study being focused on conditions below the landing minima.

---

[16] If the number of oktas sums to more than 8, a value of 8 is inserted

[17] These values are recorded when NSC, SKC and NCD are reported – as discussed in ‘Extraction of cloud groups’ on page 16

Process 3.c.i states that no ceiling is reported when the cumulative sum of oktas is zero. These values only arise when NSC, SKC and NCD are reported or forecast. Values of ceiling are nominated for these conditions, and are discussed in detail in the section below.

Approximation of weather ceiling and visibility using weather condition flags

It was relatively common for weather flags to be used in lieu of explicit cloud values, generally when conditions were reported or forecast to be favourable. Where weather forecast and observation flags (e.g. CAVOK) were present in lieu of cloud, weather phenomena and visibility information, worst-case conditions for a given flag were applied as shown in Table 17 above.

Values for all condition flags were based on the definition for CAVOK provided in AIP GEN 3.5 section 12.13. Among other things, this stipulates:

- that visibility is forecast or observed above 10 kilometres (AIP GEN 3.5 clause 12.13.1(a)
- that there is Nil Significant cloud - defining this as no cloud below the greater of 5000 feet, or the highest 25 nautical mile minimum sector altitude.

CAVOK was the only weather flag used to approximate the visibility. However, the approximation of the weather ceiling used the fields SKC, NSC and NCD in addition to CAVOK as shown in Table 17. These fields were based on the definition of CAVOK to approximate the weather ceiling as follows:

- NSC – Nil significant cloud: The definition of CAVOK contains the phrase ‘Nil significant cloud’, as described above. On this basis, the same value 4999 was used for the ceiling when NSC was reported.
- NCD – Nil cloud detected: On the basis that NSC and NCD are parrallel definitions depending on whether sourced form a human observer or an automatic weather station, the same value 4999 was used as for NSC above.
- SKC – Sky clear: Due to NSC and SKC being used in the same reports and forecasts to indicate the absence of cloud, it was assumed that there was a difference between the two. It was assumed that sky clear represented no cloud in the sky. As an upper limit, no cloud below 10,000 feet was arbitrarily chosen for comparison, input as 9999 feet.

Values were inserted to provide a lower limit of confidence given the weather flag, while not triggering any operational requirements relating to the alternate (or landing) minima. Due to this, it was expected that these values were conservative and would not affect or be relevant to the planned analysis.

In cases where values were provided in conjunction with flags, the recorded values were used over the weather flags (also shown in Table 17 above).

Vertical visibility (METARs and SPECIs only)

Vertical visibility was used to represent the weather ceiling in METAR and SPECI weather reports in the absence of any cloud groups and weather flags. For example, for a vertical visibility reported as VV003, the ceiling was recorded as 300 feet. This data is rounded down to the nearest hundred feet to represent the worst-case scenario of the ceiling.

In the absence of explicit visibility information and weather flags, vertical visibility was also used as the value for visibility.

At the time of writing, vertical visibility was no longer used in METAR and SPECI type reports, however this data was available for use in historical reports, including those used in this study.

Ceilometer and visibility sensor data anomaly (METARs and SPECIs only)

In automated weather reports, the phrase 'CLD:SKY MAY BE OBSC' has been used to indicate when no cloud is detected and visibility measured as less than 1000 metres[18]. In the absence of identifying weather ceiling data from explicit values, weather flags or vertical visibility (as described above), the weather ceiling was determined as the slant visibility height, as defined below.

The slant visibility height was defined as the height at which the reported visibility would extend at a 3-degree angle to the ground. This assumes a 3 degree approach path angle or greater, and aims to approximate the height at which an un-illuminated runway would become visible on approach. The slant visibility height calculation was calculated using the equation $h = v \sin \theta$, as shown in Figure 9 below.

If visibility was not able to be identified using the processes 0 to 2 in Table 17, this was recorded as 999 metres according to the definition, and weather ceiling was recorded as 171 feet (the slant visibility height for 999 metres).

Figure 9: Calculation of slant visibility height from reported visibility

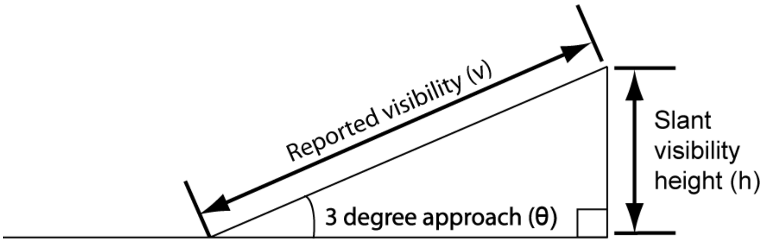

Other remarks section data

In some cases, visibility, and in particular ceiling data was present in the remarks section of reports, in particular upper level clouds. This was used in lieu of other data being present.

Average between METAR and SPECI reports

Visibility and ceiling was estimated in a very small number of cases using the average values of the reports immediately prior to and after the report with missing data. Where this approximation was made, the visibility estimate was recorded to enable further assessment and exclusion if required.

TAF and TTF values recorded for visibility or weather ceiling when the other value is determined

The validation process for TAF and TTF visibility and ceiling data was focussed on the information in each forecast being able to be interpreted by a flight crew, and therefore allowing operational decision making to take place. If a forecast was unintelligible in error, or had missing information, this was assumed to produce a safe outcome, as information would need to be sought through other means for an effective decision regarding alternate planning.

Allowances for small formatting errors were made in cases where information was able to be determined. If cloud information was not present in a forecast or able to be identified, such as when fog was forecast, the cloud base was assumed to be the least limiting factor and set to 4998 feet for identification purposes.

In a small number of cases, visibility information was omitted from forecast information. In these cases, the visibility was set to 9998 metres for identification purposes and the cloud information was assumed the most limiting factor. In cases where both visibility and cloud information was missing, these forecasts were excluded from analysis.

---

[18] Based on the definition for CLD: SKY MAY BE OBSC provided in remarks section of the Bureau of Meteorology 'METAR / SPECI' fact sheet published 11 February 2015.

Formation of relational instrument approach procedure data framework

Instrument approach data was used to define the landing and alternate minima criteria for every selected aircraft given forecast and reported weather conditions at each aerodrome. A relational data structure was selected to store the instrument approach data for efficiency and to allow the storage of each data element under a single field.

For example, an instrument approach may be listed with four separate minima for each aircraft category, which in a flat table structure requires four separate columns, and in the structure used is then converted to a single column in four separate rows. As stated in the section 'Formation of relational weather data framework' on page 19, a relational structure also provides the benefit of performing aggregate operations and calculations, for example determining the approach with the lowest possible landing minima for a specified aircraft in reported weather conditions.

The data structure developed is shown in Figure 10 below.

Instrument approach procedure extraction into a relational data structure

The flat data table (Table 37) described above was imported in delimited format into an SQL database. This was then extracted into a relational data structure of seven tables, as illustrated in the orange boxes in Figure 10 below. This diagram also shows how aircraft specification data is linked (shown in green), the four instrument approach and aircraft related lookup tables described above and relevant table links (light blue), and where waypoint and runway data (blue) is joined.

The tables labelled 'Aircraft approach minima' and 'Aircraft alternate minima runway' are tables created by linking runway, aircraft specification and instrument approach data together, and were used as the primary reference tables when linking with reported and forecast weather conditions. Combining this data is discussed in 'Instrument approach selection for each METAR and SPECI report using landing minima criteria' on page 58.

The tables labelled 'approach minima to alternate minima' and 'Aerodrome approach minima to circling minima' were used to link the approach minima table to both the aerodrome alternate minima and aerodrome circling minima tables. These links were created to ensure that the alternate or circling minima data was explicitly linked to the minima criteria in each approach plate.

This data was generally the same within each aircraft performance category (but not assumed to be) for all approaches at an aerodrome. For this reason, the data was stored under the location, rather than being duplicated for every approach plate. This was also designed to allow the investigation of variation between different runways and approach plates in different reported and forecast weather conditions.

Figure 10: Instrument approach and aircraft specification data structure

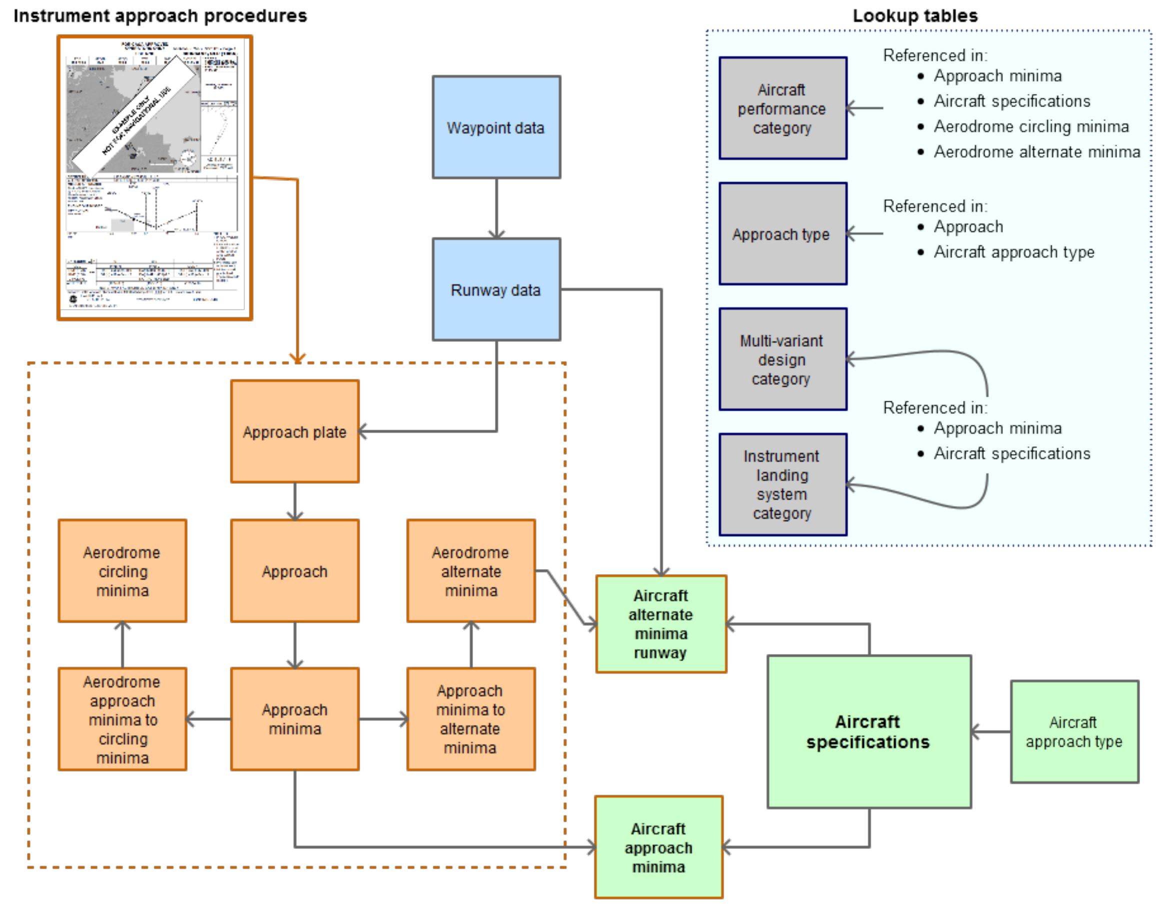

Data was compiled in the following sequence:

1. Lookup tables created (Figure 10 Light blue boxes)
2. Aircraft specifications imported (Figure 10 Green box)
3. Waypoint data imported from ATSB database (Figure 10 Blue box)
4. Form runway data table (Figure 10 Blue box)
    a. Runway data imported and copied from Airservices Australia data
    b. Runway directions (to the nearest degree) appended to runways of interest
5. Create instrument approach procedure framework (Figure 10 Orange boxes)
6. Import instrument approach procedure flat file data (Figure 10 Orange boxes)
    a. Populate approach plate table from instrument approach procedures
        i. Create foreign key link to runway data table
    b. Populate approach table data for each approach plate entered above
        i. Foreign key links to approach plate and approach type
    c. Populate approach minima table
        i. Separate aircraft category information (Minimum altitudes and heights, and minimum visibility) into separate rows (shown in table Table 37 on page 101 as columns suffixed by category eg. MinimumVisibility_A)
        ii. Duplicate linking fields across categories (eg. ILS and MVD Id columns – non-suffixed column names shown in Table 37 on page 101)
        iii. Foreign key links: Aircraft performance category, Approach, ILS Category, MVD category
    d. Populate alternate minima data tables

        i. Separate aircraft category information, including special alternate minima data (as per step c.i above)
        ii. Populate aerodrome alternate minima table
            1. Links to Aircraft performance category and Waypoint table
        iii. Populate linking table to link approach minima data to the alternate minima data that was published on the same approach plate
    e. Populate circling minima data tables
        i. Separate aircraft category information data (as per step c.i above)
        ii. Populate aerodrome circling minima table
            1. Links to Aircraft performance category and Waypoint table
        iii. Populate linking table to link approach minima data to the circling minima data that was published on the same approach plate

Reference Frames

To align reference frames, all approach procedure variables were modified to align with weather forecasts and observations unless otherwise specified. All directions from the aerodrome were converted to a true north reference frame, and all heights were converted to above the aerodrome reference point to align with the cloud ceiling reference and wind and weather directions in TAF, METAR and TTF weather products. This includes precision approach heights above runway thresholds, which are also converted to height above the aerodrome reference point.

Evaluation of weather reports and forecasts

The following section describes the process of combining instrument approach procedures, aircraft specifications and weather reports and forecasts together. This was performed to determine if the alternate and landing minima were breached for every forecast and report for the aircraft and locations selected. The alternate minima assessed forecasts were then evaluated against the landing minima assessed weather reports for over-lapping times.

The evaluation of weather reports and forecasts was performed using formed data as described in the section ‘Relational data structure formation, extraction and validation’ above. The evaluation process followed several key processes, as depicted in Figure 1 on page 8, which also follows the following main headings in this section.

- Instrument approaches linked to aircraft models
- Determining possible alternate minima applicable to an aircraft
- Deletion of selected records for re-calculation
- Autonomous exclusion of invalid forecasts based on misaligned timing segments
- Timing buffers added to TAFs
- Calculation of TTF segment effective start and end times
- Instrument approach selection for each METAR and SPECI report using landing minima criteria
- Assessment of TAFs and TTFs against alternate minima criteria
- Comparison of overlapping TAF and TTF forecasts with METAR and SPECI reports
- Overall forecast assessments for operational reliability.

Instrument approaches linked to aircraft models

Visibility and cloud base requirements within instrument approach procedures were used as the basis for determining if conditions were reported above or below the landing minima, and additionally, if conditions were forecast above or below the alternate minima. To study the effect that this would have on a nominated aircraft model, instrument approaches were selected that a nominated aircraft was equipped to fly.

The subset of approaches linked to each nominated aircraft were stored in a linking table ‘Aircraft approach minima’ as shown in Figure 10 and Figure 11. Approaches from this table were selected based on an algorithm to select the most appropriate runway for landing, which assessed the wind speed and direction, and the accuracy of each approach. This algorithm is discussed in depth in ‘Instrument approach selection for each METAR and SPECI report using landing minima criteria’ on page 58.

At the time of writing, the Boeing 737-800 aircraft was used as the basis for all approach selections. The B737-800 was assumed to be equipped with instrumentation suitable to fly all types of published approaches (as shown in Table 47 on page 106). However, the selection of instrument approaches was restricted within the published performance criteria set forward within each approach procedure.

Information regarding the determination of each of these elements for aircraft and instrument approach procedures is contained in sections ‘Formation of instrument approach procedure flat data file’, ‘Compilation of aircraft specification flat data file’, and ‘Instrument approach procedure extraction into a relational data structure’ on pages 14, 16 and 46.

Figure 11 and Table 19 below show the processes used to match instrument approach procedures with aircraft models. Figure 10 and Figure 11 show the overall linking table labelled ‘Aircraft Approach Type’, which contained one row for every aircraft model and approach that each aircraft was expected to be equipped to fly. The four upper rectangles represent data from the tables created in an earlier process, and are also shown in Figure 10 above. The data matching process followed these stages:

- A link was created between every instrument approach and aircraft (from table ‘aircraft approach type’) using the approach type category (Table 47), and formed into a temporary table (labelled as ‘Approach types’ in Figure 11)
- Data was removed from the temporary table if the aircraft performance specifications did not meet five requirements published in the instrument approach procedure, based on:
- Aircraft performance category
- Aircraft less than or equal to instrument approach category
- Multi-variant design (MVD) category
- Aircraft less than or equal to instrument approach category
- ILS capability
- Aircraft capability less than or equal to instrument approach infrastructure
- Required navigational performance (RNP)
- Aircraft capability more accurate or equal to required performance
- Minimum missed approach climb gradient
- Aircraft performance greater than minimum climb gradient
- Data meeting all criteria was inserted into the table Aircraft approach minima, as shown in Figure 10 and Figure 11

The last four of the five requirements were only applicable to selected approaches. If a criterion was not applicable, it was ignored as a non-limiting factor for linking the approach to the aircraft.

Figure 11: Matching instrument approach procedures to aircraft models

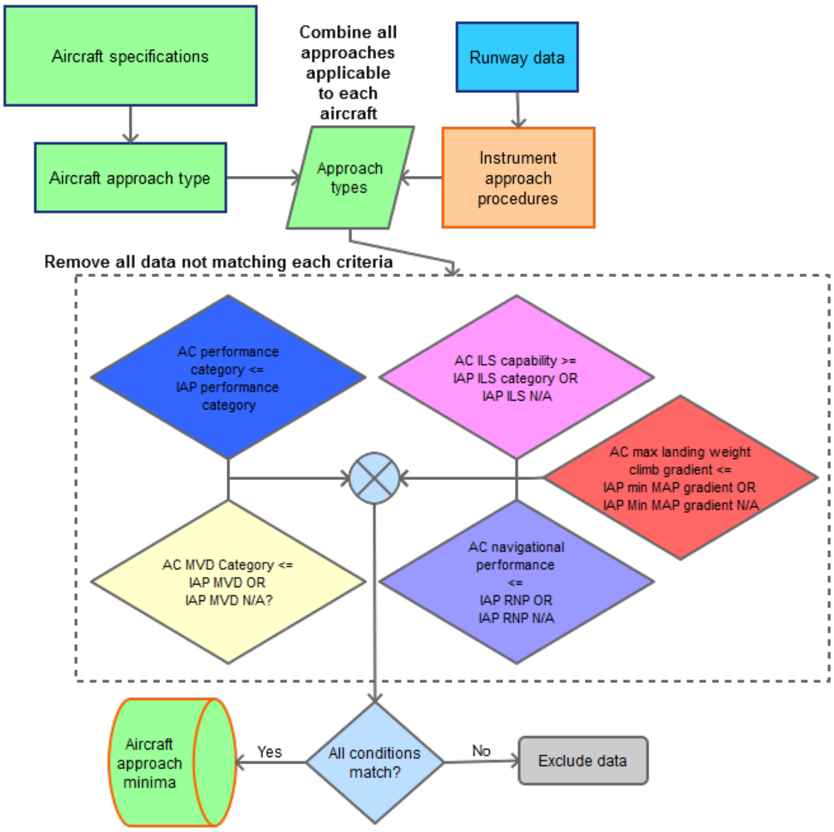

Table 19: Sample instrument approach procedures matched and excluded for aircraft based on specifications

ICAO Code	Aircraft	Approach	Performance Category		MVD Category		ILS Category		RNP		Min MAP Grad	
		Aircraft (A) / IAP	A	IAP	A	IAP	A	IAP	A	IAP	A	IAP
YPAD	B737-800	S-I VOR/DME RWY 05	D	D	Narrow body	NULL	IIIb	N/A	0.05	NULL	7.8	3.6
		RNAV-U RNP RWY 23		D		Regional		N/A		0.3		5.6
		S-I ILS RWY 23		D		NULL		I		NULL		NULL
		RNAV-V RNP RWY 23		D		Wide body 2 engine		N/A		0.1		5.6
Sample	Sample	S-I ILS RWY 25	D	D	Wide body 4 engine	NULL	I	II	0.25	NULL	4.0	5.6
		RNAV-U RNP RWY 07		B		Narrow body		N/A		0.1		NULL

Table 19 shows a sample set of data to illustrate the output following the matching of approach types as shown in Figure 11, and described above. The colours on the top row correspond to each of the criterion listed in Figure 10. The sub-columns labelled ‘A’, and ‘IAP’ represent that the underlying data is from aircraft specifications and instrument approach procedures respectively, and are also colour coded with respective sources in Figure 11.

Any approach with light red shading under any criterion indicates where an aircraft specification has not met the requirements of the instrument approach procedure. If any of the criterion did not meet the requirements, they were not included in the linking table ‘Aircraft approach minima’ shown in Figure 10. For example, the fictitious RNAV-U RNP RWY 23 approach would not be linked to the Boeing B737-800 aircraft due to the requirement of a regional multi-variant design category, and the B737 being in the larger narrow body category. The example in Table 19 shows three approaches that would be included in the linking table (shaded light green), and three approaches that would be excluded (shown in light red).

Determining possible alternate minima applicable to an aircraft

Following the creation of the link between aircraft specifications and instrument approach procedures as described above, another table was created to link published alternate minima visibility and ceiling limits with aircraft specifications. The table created was named ‘Aircraft alternate minima runway’ after the three main elements being explicitly linked in the process. The diagram of data flow for this table is shown in Figure 10 on page 47.

This table was formed to allow the display of all alternate minima combinations for each runway and aircraft model. The selection of a specific alternate minima value given forecast conditions is discussed in section ‘Assessment of TAFs and TTFs against alternate minima criteria’ on page 67.

This table is formed from alternate minima information contained in each of the instrument approaches found to apply to nominated aircraft as defined by the process above. For every applicable linked instrument approach procedure and aircraft combination, one alternate minima value was inserted for the aircraft performance category.

For some instrument approaches, the special alternate minima was also available for the selected aircraft. In these cases, an additional row of data was added with the special alternate minima values.

Due to alternate minima generally being published as the same value within each aircraft performance category at an aerodrome, almost all of the values in the list were the same. Because of this, these values were condensed to a list of distinctive rows for each unique combination of runway, aircraft and alternate minima values, and are illustrated in Table 20. For example the ILS and RNAV-V approaches for runway 23 contain the same alternate minima information regarding performance category D aircraft, and as such, this was stored as a single row.

Variation in alternate minima values between runways are accounted for by only condensing rows with identical visibility and ceiling values. For example, Table 20 shows the runway 05 approaches applying to Category D aircraft. If no special alternate minima was listed on this approach plate, or any others for runway 05, the special alternate minima values would not apply to runway 05, and would not have been recorded in the table (however in reality this was not the case for Adelaide).

This accounts for effects such as the selection of an alternate minima based on an expected approach runway from a forecast, and conditions favouring a different runway upon arrival. It also allows for unexpected discrepancies in alternate minima data between different approach plates.

Table 20: Sample formation of data rows for unique alternate minima combinations for specified aircraft

Source tables						Target table			
Aircraft approach minima			**Approach minima to alternate minima**			**Aircraft alternate minima runway**			
Aircraft	Approach	Runway	Category	Visibility	Ceiling	Runway	Aircraft	Visibility	Ceiling
B737-800	S-I ILS RWY 23	23	D	7000	1313	23	B737-8	7000	1313
B737-800	RNAV-V RNP RWY 23	23	D	7000	1313				
B737-800	S-I ILS RWY 23	23	Special	4000	850	23	B737-8	4000	850
B737-800	VOR RWY 05	05	D	7000	1313	05	B737-8	7000	1313

Deletion of selected records for re-calculation

To allow the re-calculation of data following a change to an algorithm at any of the main stages of the analysis, algorithms were created to delete data. This allowed existing, complete data sets to remain unaffected. At the time of writing, deletion of records by location could be performed if new data became available, or changes to algorithms were performed affecting some locations and not others.

Due to the referential integrity of the database created, items could not be deleted if these left any child items in the data set. For this reason, it was considered very unlikely that the unintended analysis of data thought to be removed could occur.

Autonomous exclusion of invalid forecasts based on misaligned timing segments

To evaluate if the start and end time of some TAF and TTF segments were in the correct sense, and that concurrent permanent segments were in alignment, two checks were performed to exclude invalid data. A description of the what is meant by ‘Permanent’ segments are described in ‘TAF permanent change segment adjustments’ on page 54.

This involved reviewing the TAF body, FM and BECMG segment timing for TAFs and the TTF body and FM segments for TTFs. The checks were as follows:

- Segment start time greater than or equal to the next segment start time
- Segment start time greater than the segment end time

Any TAFs or TTFs identified were excluded from any further analysis. Further discussion on the validation of data is contained in ‘Validation and summary of operational reliability of each TAF and TTF assessment’ on page 81.

Timing buffers added to TAFs

At the time of writing, operational requirements[19] stipulated that for the purposes of planning an alternate, conditions forecast by a TAF that trigger an operational requirement be treated as starting and ending 30 minutes prior to and after the predicted time, with the exception of PROB forecast periods. These time adjustments may require actions such as taking on additional holding fuel, or planning with an alternate aerodrome. To assess if observed conditions were as predicted in each

---

[19] This was stipulated in AIP ENR 1.1 subsection 58, and is summarised on page 57.

forecast, the effective time (the time taking into account the required buffer) of each predicted change was required to be determined. This data was inserted into a new table of TAF sub-segments, and is described as follows.

To account for the requirements to adjust time, additional time fields were appended to every forecast segment where applicable. Figure 12 and Table 21 illustrate the processes used to adjust the start and end times of each TAF segment. The stipulated contingency time period, labelled ‘t_c’ in Figure 12 represents the time offset to be applied surrounding an operational requirement. As mentioned above, the time offset applied was 30 minutes, which could be adjusted to allow the assessment of a different offset in future analyses.

Figure 12: Formation of TAF sub-segments accounting for start and end time adjustment requirements

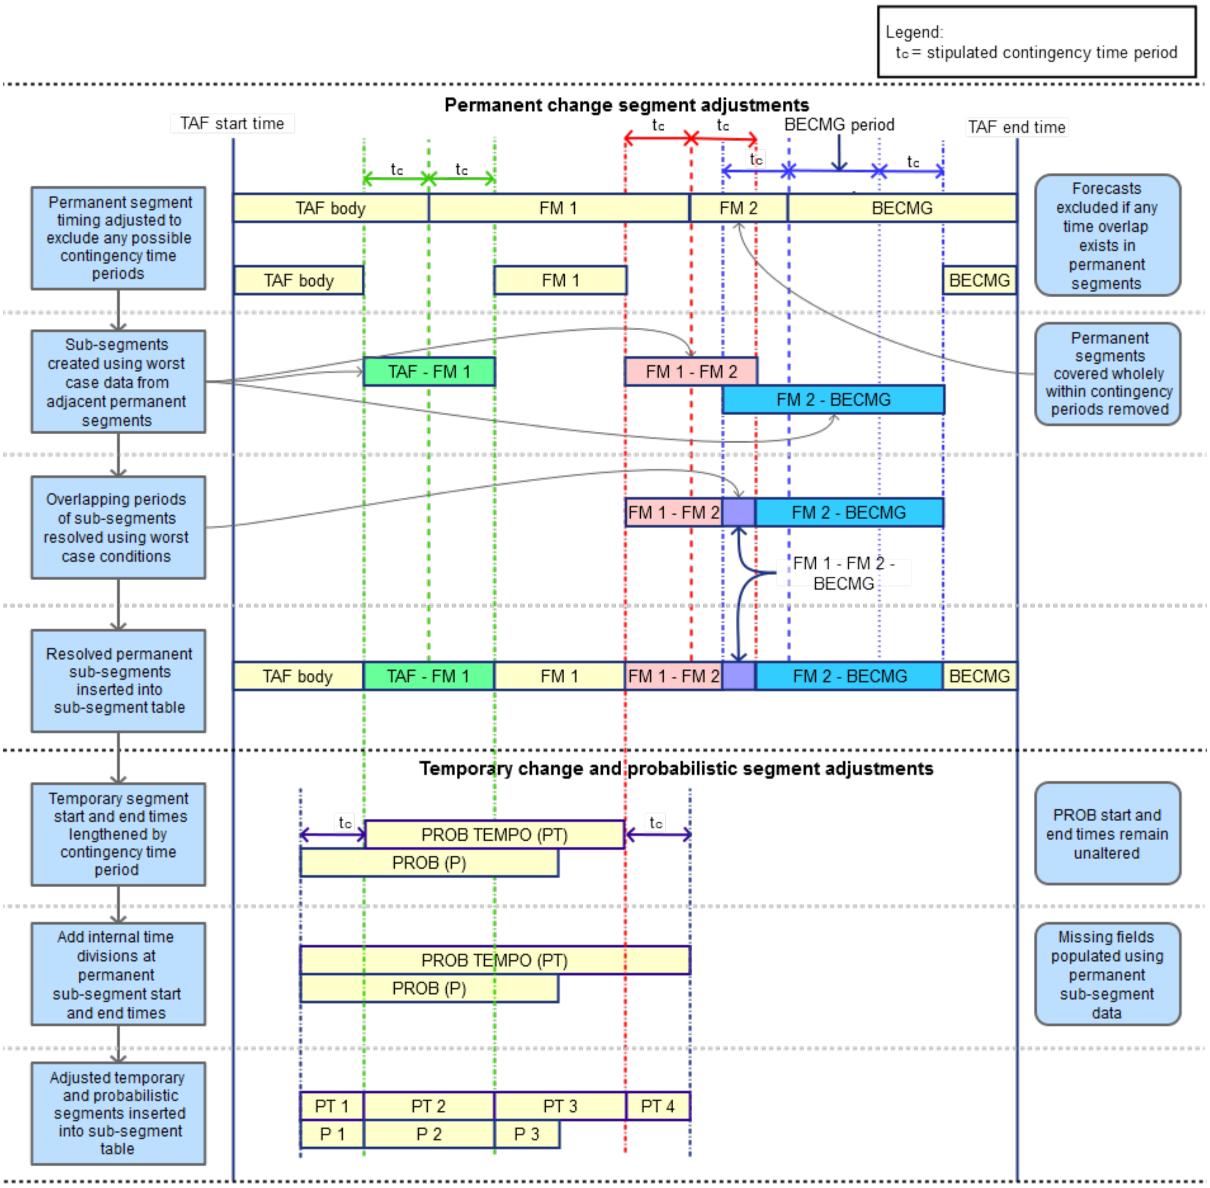

Table 21: Example adjustment of stipulated 30 minute start and end time buffers in permanent TAF segments and creation of sub-segment times

Original TAF decode table				**New TAF sub-segment table**		
Segment	**Segment Start Time**	**Segment End Time**	**BECMG Period End**	**Segment**	**Sub-segment start time**	**Sub-Segment end time**
TAF	00:00	14:00		TAF	00:00	13:30
				TAF – FM 1	13:30	14:30
FM 1	14:00	16:00		FM 1	14:30	15:30
				FM 1 – FM 2	15:30	16:15
FM 2	16:00	16:45				
				FM 1 – FM 2 – BECMG	16:15	16:30
				FM 2 – BECMG	16:30	19:00
BECMG	16:45	00:00	18:30	BECMG	19:00	00:00
PROB (P)	13:00	15:00		P 1	13:00	13:30
				P 2	13:30	14:30
				P 3	14:30	15:00
PROB TEMPO (PT)	13:30	15:30		PT 1	13:00	13:30
				PT 2	13:30	14:30
				PT 3	14:30	15:30
				PT 4	15:30	16:00

TAF forecast segments were divided into two groupings for assessment:

- permanent significant changes (TAF body, FM and BECMG)[20]
- temporary variations and probabilistic changes (INTER, TEMPO, PROB, PROB INTER, PROB TEMPO)[21]

TAF permanent change segment adjustments

The upper portion of Figure 12 shows the main processes used to adjust the timing and data of each permanent TAF segment, as stipulated in the AIP requirements. This involved the creation of new ‘worst-case’ combined segments in between adjacent forecast segments and the retention of the original forecast data not falling within the stipulated time adjustment periods. An example of this process using the same labels and colours in Figure 12 is shown in Table 21 above.

The creation of new, ‘worst-case’ combined segments using data from adjacent forecast segments was based on a rationale of extending the intention of the stipulated requirements. That is, when a forecast predicts conditions creating an operational requirement, an extension of the time at both the beginning and end of the predicted period is required.

It was assumed that including only the most severe value from adjacent forecasts for each parameter would provide an equivalent result to overlapping the data for the purpose of the assessment of the alternate minima, as discussed on page 67. This is because the alternate minima criterion were independent of one another and other parameters, meaning that operational requirements are triggered if any of the criterion were met. In addition, the use of ‘worst-case’ data was assumed to be conservative, and to follow the intent of the stipulated requirements.

---

[20] Based on Aeronautical Information Publication ENR 1.1 subsections 58.2.3 and 58.2.7 dated 20 August 2015.

[21] Based on Aeronautical Information Publication ENR 1.1 subsection 58.2.8, dated 20 August 2015.

The start of the process involved ‘quarantining’ periods of time not potentially affected by time adjustment requirements. These are shown in Figure 12 and Table 21 as the light yellow shortened segments ‘TAF body’, ‘FM 1’, and ‘BECMG’. For example, the start and end of the segment labelled FM 1 is shortened by 30 minutes at each end from 1400 and 1600 to 1430 and 1530.

Sub-segments were then created in the time gap formed in the previous step. In the previous example, a segment labelled ‘TAF – FM 1’ was created between 1330 and 1430.

In this example, the ‘FM 2’ segment is removed because it falls completely within the contingency time requirements, however, it is accounted for in the combined data described below.

The end of a ‘becoming-period’ (BECMG) (in contrast to the end of the BECMG segment) was used as the reference to adjust the end of the contingency time at the interface between the end of the previous segment and beginning of a BECMG segment. This is shown in Figure 12 and Table 21, where the ‘FM 2 – BECMG’ sub-segment extends 30 minutes beyond the BECMG period end of 18:30.

In cases where an overlap exists between created combined segments (when a forecast segment was less than twice the contingency time, ‘t_c’), a further process was implemented to coalesce overlapping data. This followed the same approach of using the worst-case data from each overlapping combined segment. A combined segment is depicted in purple in Figure 12 and Table 21, and is labelled ‘FM1 – FM2 – BECMG’, containing the resultant worst-case data in each of these three segments. This section of the algorithm also accounts for multiple overlaps if many relatively short segments are adjacent to one another. Segments with no time (0 minutes) are accounted for in this algorithm, with the worst case buffers applying around this transition point. However, checks were implemented to prevent segments with a published start time after the end time being processed.

At the end of the permanent sub-segment creation process all combined segments were non-overlapping and adjacent to the original, time adjusted permanent forecast segments. Any segment transitions occurring within the contingency time period of the start and end of the TAF validity were ‘trimmed’ to ensure that no overlap beyond the validity period existed. A ‘check-sum’ was performed adding the durations of all TAF permanent sub-segments and comparing this with the duration of the TAF validity period. A message was produced in the computer program to confirm that all check-sums were valid, or if any anomalies were identified.

Forming non-overlapping permanent time elements was used to form an operational ‘baseline’ assessment for every minute of a TAF in comparison to weather reports, followed by the assessment of whether any temporary segments changed the ‘baseline’ operational requirements. This is discussed on page 77.

The new TAF ‘sub-segments’ were then populated with the worst-case data from the adjacent, original permanent segments. This included visibility, ceiling, wind speed, and weather phenomena, including thunderstorms. An example output of this process is shown in Table 22, using the same segments shown in Figure 12 and Table 21 above. The figures in **bold** in the original TAF segments indicate the source of data in the new sub-segments surrounded by a shaded box.

Table 22: Example of worst-case selection for permanent TAF sub-segments

Segment	Visibility	Effective ceiling	Wind direction	Average wind speed	Thunderstorms	Showers	Rain	Fog
TAF	9999	4999	**220**	**17**				
TAF – FM 1	9000	1200	220	17				
FM 1	**9000**	**1200**	**260**	**14**				
FM 1 – FM 2	9000	1200	260	14		1	1	
FM 2[22]	*9999*	*11000*	*340*	*14*		***1***	***1***	
FM 1 – FM 2 – BECMG	1500	1200	210	16		1	1	1
FM 2 – BECMG	1500	3000	210	16		1	1	1
BECMG	**1500**	**3000**	**210**	**16**				**1**

The most severe weather parameter value for visibility and effective ceiling from each adjacent segment was inserted into the new TAF sub-segments. The entire wind group field (including the direction, average and maximum) was extracted from the adjacent TAF segment with the largest average wind speed. If the same comparison value for visibility, ceiling or average wind speed was present in both adjacent segments, the earlier of the two segments were used.

Table 22 also shows data from the removed segment, labelled ‘FM 2’ being used to populate sub-segments where a time overlap applies and the original ‘FM 2’ segment contains any worst case element. In this case showers and rain were predicted in the segment, which is represented in all segments falling within the ‘FM 2’ contingency overlap as shown.

TAF temporary and probabilistic change segment adjustments

The original start and end times of any temporary TAF segments (TEMPO, INTER, PROB INTER and PROB TEMPO) were extended by the contingency time of 30 minutes. These were inserted into the TAF sub-segment table mentioned above with adjusted start and end times. These segments were treated as being independent of one another. No changes were made to the timing of information in probabilistic segments ‘PROB’ without any forecast temporary time period. These segments were also inserted into the TAF sub-segment table. This is shown in Figure 12 and Table 21.

In many cases, wind, visibility and cloud base forecast parameters were not published explicitly in all of these segments. For example, when fog was forecast, only visibility was published in that part of the forecast, with wind and cloud predictions omitted.

It was assumed that for wind, visibility and cloud base information not explicitly published in temporary or PROB segments, this information could be taken from the permanent segment data forecast to occur at the same time. This is because it was expected that operational personnel would look to ‘fill in the gaps’ when interpreting a forecast when trying to assess the predicted conditions that may affect a safe landing. Having this information complete for these sections also assisted in the comparison of these segments with observations, as described in ‘Comparisons between forecasts and reports’ on page 73.

---

[22] In this example, the segment labelled FM 2 is removed due to the contingency time overlap, however is listed to illustrate the data source for the created sub-segments.

To achieve this, internal timing divisions were made in every extended temporary and probabilistic sub-segment using the same times established for the permanent sub-segments. The addition of these time breaks is shown in Figure 12, with an example showing timing in Table 21. The internal time divisions are always aligned between all temporary and PROB segments within each TAF because they are defined using the non-overlapping TAF sub-segments created. In this example, it can be seen that the internal timing divisions are aligned for both the PROB and PROB TEMPO segments to the TAF, TAF – FM 1, FM 1, and FM 1 – FM 2 segments. Note also that the PROB segment is not extended, as this was not stipulated in the AIP.

In cases where visibility, cloud base or wind information was missing from temporary or PROB segments, it was filled using the aligned permanent sub-segment information forecast to occur at the same time within each time division.

Table 23 continues the example from Figure 12 and Table 21. This shows that the missing wind information shown in *italics* for the original PROB segment was inserted from the three permanent sub-segments forecast to occur at the same time. Values contained in the original PROB segment are shown in **bold**. Note also that the effective ceiling value for the PROB segment labelled 4998 indicates that this was artificially inserted in the data as a least limiting factor, as discussed in the section 'TAF and TTF values recorded for visibility or weather ceiling when the other value is determined' on page 45.

In contrast to the PROB segment, the wind, visibility and cloud base information was contained in the PROB TEMPO segment in Table 23. Therefore, this only uses values contained from the original forecast segment.

The remaining information such as forecast weather phenomena (Fog, Thunderstorms and Rain in the example in Table 23) and characteristics were unchanged in the data, and were not combined with any other segments, even where they were overlapping.

Table 23: Determination of temporary and probabilistic TAF sub-segments complemented with permanent TAF sub-segment data

Original temporary and probabilistic segments					
Segment	**Parameter**	**Value**	**Segment**	**Parameter**	**Value**
PROB	Visibility	**0500**	PROB TEMPO	Visibility	**2000**
	Ceiling	**4998**		Ceiling	**1000**
	Wind direction			Wind direction	**VRB**
	Average wind speed			Average wind speed	**20**
	Phenomena	**FG**		Phenomena	**TS** **RA**

Permanent sub-segments			New temporary and probabilistic segments					
Segment	**Parameter**	**Value**						
TAF	Visibility	*9999*	P 1	Visibility	**500**	PT 1, PT 2, PT 3, PT 4		
	Ceiling	*4999*		Ceiling	**4998**			
	Wind direction	*220*		Wind direction	*220*			
	Wind Speed	*17*		Average wind speed	*17*			
TAF – FM 1	Visibility	*9000*	P 2	Visibility	**500**		Visibility	**2000**
	Ceiling	*1200*		Ceiling	**4998**		Ceiling	**1000**
	Wind direction	*220*		Wind direction	*220*		Wind direction	**VRB**
	Wind Speed	*17*		Average wind speed	*17*		Average wind speed	**20**
FM 1	Visibility	*9000*	P 3	Visibility	**500**			
	Ceiling	*1200*		Ceiling	**4998**			
	Wind direction	*260*		Wind direction	*260*			
	Wind Speed	*14*		Average wind speed	*14*			

No timing adjustments performed for TTFs

TTFs had no timing adjustments performed.

This is in accordance with the current Australian flight planning rules, and was implemented to simulate the system in its current state, and to establish a baseline for further analysis.[23]

Instrument approach selection for each METAR and SPECI report using landing minima criteria

Establishing when METARs and SPECIs reported conditions below the landing minima criteria was used as the primary criterion for assessment of TAF and TTF forecasting, and is discussed in detail in 'Assessment of METARs and SPECIs below the landing minima' on page 65. This section describes how one aircraft-specific instrument approach (described on page 48) was selected for every METAR and SPECI report, including the assessment of whether the reported landing conditions were above or below any of the landing minima thresholds.

---

[23] AIP ENR 59 dated 18 Aug 2016.

A number of comparisons were performed between each METAR and SPECI, and the list of possible aircraft and instrument approach combinations. Conditions suitable for landing were assessed against the requirements put forward in the AIP documentation and instrument approach procedures applicable to the conditions and aircraft.

Selection and calculation of landing minima criterion data fields

The columns in Table 26 under the label 'METAR or SPECI comparison with instrument approach and aircraft data' show comparisons between published landing minima thresholds (sourced from instrument approaches (labelled IAP) and aircraft (labelled AC)) and METAR and SPECI reports (labelled M).

Some of these values required selection (where more than one value was available) and calculation (such as wind component values) prior to comparison as follows.

Selection and calculation of approach visibility and runway visual range

Depending on the type of approach, every instrument approach procedure contains one or both of the landing visibility minimum and a minimum runway visual (RVR) range. Due to only one value for each visibility and RVR as the landing minima thresholds, algorithms were required to determine a single value for each parameter sourced from METAR and SPECI reports.

Approach visibility

To account for cases where the minimum visibility and minimum visibility direction was reported in addition to a prevailing visibility, an algorithm was developed to approximate the reported visibility for each runway approach. This algorithm used the existing data field formed for visibility as described on page 41 as the assumed prevailing visibility and the reported minimum visibility.

The resultant calculation is a runway specific visibility value, referred to as the approach visibility, and represents an approximation of the visibility that a flight crew may experience on approach to each runway. If no minimum visibility was reported, the approach visibility was reported as the prevailing visibility.

When the directional minimum visibility (assumed to be relative to true north) was reported, a sinusoidal approximation was used to determine the likely visibility when approaching the aerodrome from the hemisphere of the reported reduced visibility. For example if the reduced visibility is reported from the north-east, it is assumed that the visibility is reduced for any approach flown toward the aerodrome from the north-west through north-east to the south-east. This is shown in Figure 13.

In this example, assume that the visibility group in the report is 9000 2000NE. The angle β representing the minimum visibility direction of NE is recorded as 45 degrees, and ϕ, the angle representing the runway (and approach) direction is 270 degrees. As was the case with all magnetic runway directions, these were converted to a true north datum prior to comparison.

It was assumed that the visibility would be reduced in the direction of the reported minimum visibility, and any aircraft flying *from* that direction would be affected, the opposite direction to the approach. Therefore, the reciprocal runway direction (the opposite runway) was used as the comparison angle, shown in the equations as recip(ϕ). The absolute value of the difference ($\Delta\theta$) between these angles was determined, to establish the difference in radial position from the aircraft to the reported minimum visibility. In this example, the difference would be: $\Delta\theta = \text{abs}(45 - (360 - 270)) = 45$ degrees. As the difference was less than 90 degrees, it is assumed that the aircraft would be affected on approach by the reported minimum visibility.

A sinusoidal model was selected to represent the decaying visibility due to the relatively stronger weighting of radial directions near the stated minimum visibility and rapidly decreasing weighting away from this value. This was desired due to the prevailing visibility being at least half of the horizon circle[24]. Table 24 shows a number of outputs from the calculation of approach visibility by varying runway directions and holding the observed conditions constant. This shows the variation between the minimum and prevailing visibility. For example, the reciprocal runway directions 135 and 315 have the same value of 9000 metres visibility, decreasing to halfway between the prevailing and minimum visibility of 5500 metres when the relative angle drops to 30 degrees, down to the minimum visibility of 2000 metres when the relative angles coincide.

Figure 13: Approximation of approach visibility from METAR and SPECI minimum and prevailing visibility fields

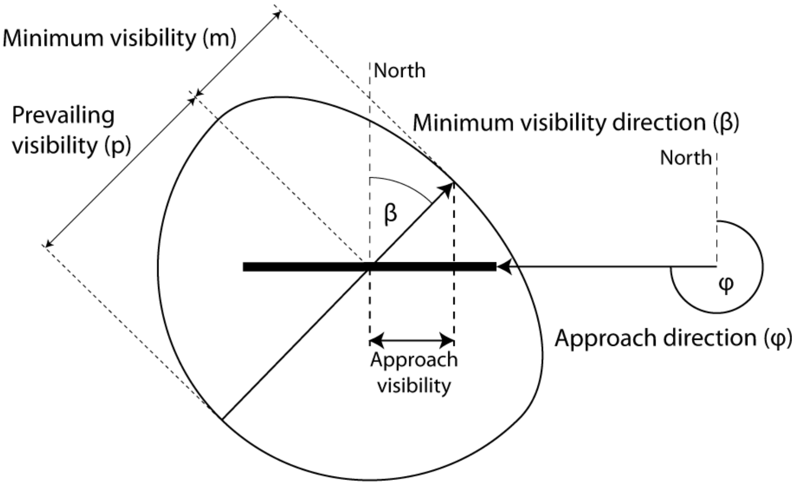

$$Let\ recip(\varphi) = \varphi + 180\ when\ \varphi < 180; recip(\varphi) = \varphi - 180\ when\ \varphi \geq 180$$
$$\Delta\theta = abs(\beta - recip(\varphi)) for\ abs(\beta - recip(\varphi)) < 180$$
$$or\ \Delta\theta = 360 - abs(\beta - recip(\varphi)) for\ abs(\beta - recip(\varphi)) \geq 180$$

$$If\ \Delta\theta \leq 90, \qquad Approach\ visibility = m + (p - m) \times sin\left(\Delta\theta \times \frac{\pi}{180}\right)$$

$$If\ \Delta\theta > 90, \qquad Approach\ visibility = p$$

Table 24: Sample output from approach visibility function

Prevailing visibility (p)	Minimum visibility (m)	Minimum visibility direction (β)	Runway direction (φ)	Recip(φ)	Δθ	Approach visibility
9000	2000	045	315	135	90	9000
			285	105	60	8062
			270	90	45	6949
			255	75	30	5500
			140	320	85	8973
			240	60	15	3811
			135	315	90	9000
			225	45	0	2000

Selection and calculation of runway visual range

Runway visual range data was extracted for published instrument approaches where this was included. At the time of writing, RVR only applied to instrument landing system (ILS) approaches. These lower thresholds were used where stipulated instead of visibility for the assessment of the

---

[24] Australian Bureau of Meteorology METAR/SPECI brochure 11 February 2015

landing minima, as described in ‘Assessment of METARs and SPECIs below the landing minima’ on page 65.

RVR information in SPECIs (and a small number of METARs) contained one or two RVR figures depending on observation of significant fluctuations in the previous 10 minutes. A single nominal value of was determined for each RVR group linked to allow comparison with relevant instrument approach procedures. This information was assumed to represent the average runway visual range during the period of the report.

Note that the rationale behind the selection of RVR is to determine the closest expected value, not the most conservative value. This was because it was desired to estimate actual conditions.

Due to RVR being specific to runways, multiple RVR groups were possible (and common) in each METAR or SPECI. Figure 14 shows the processes used to nominate a runway visual range value for a given RVR group, linked to a runway, with Table 25 showing sample outputs for each of the possibilities from this algorithm.

Figure 14: Determination of effective runway visual range

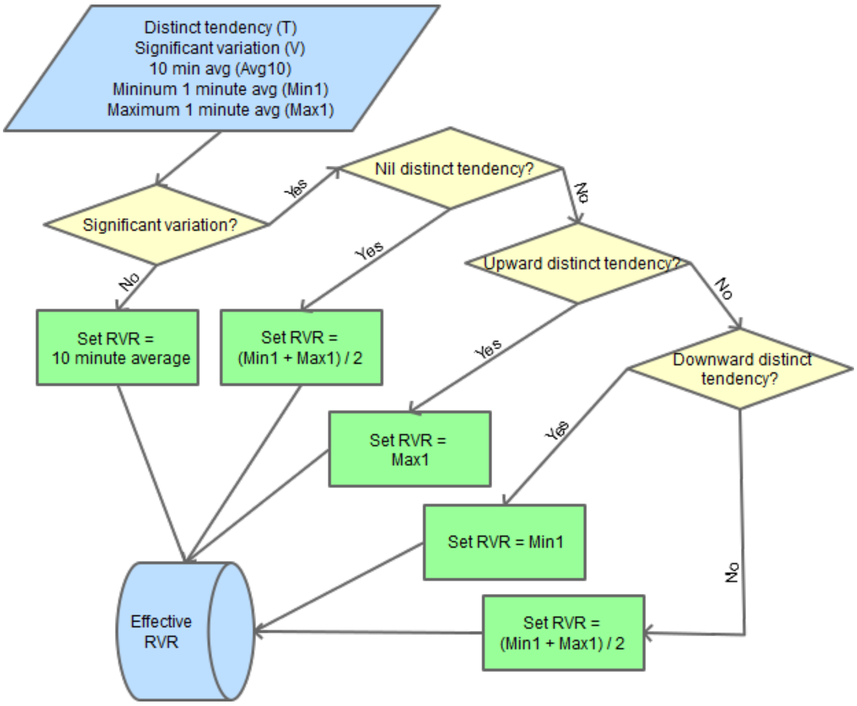

In cases where no significant variation was reported, the 10-minute average was the only figure available for RVR. Where there was significant variation and a distinct upward or downward trend observed, the minimum or maximum 1-minute average value was used corresponding to this trend. For example, Table 25 shows an upward trend with minimum and maximum 1-minute averages of 450 and 900 metres respectively. In this case, 900 metres was selected as it was assumed that the upward trend would apply.

If no trend existed, the average of the minimum and maximum 1-minute values was nominated as the effective RVR for the runway. When significant variation was reported, but no distinct tendency information was present in the data, the average of the minimum and maximum 1-minute values were used.

Table 25: Sample output from effective runway visual range algorithm

Original data	Extract					Output
Runway visual range text	**Significant variation**	**Distinct tendency**	**10 minute average**	**Minimum 1 minute average**	**Maximum 1 minute average**	**Effective runway visual range**
R27/P2000N	No	Nil	2000	N/A	N/A	2000
R27/0450V0900U	Yes	Upward	N/A	0450	0900	900
R16/0125V0450D	Yes	Downward	N/A	0125	0450	125
R27/0325V0400N	Yes	Nil	N/A	0325	0400	363

Alignment of ceiling data

Prior to comparison of cloud base (ceiling) between instrument approach procedures and weather reports, ceiling values in instrument approach data was adjusted for elevation. This was performed by subtracting the aerodrome elevation from the minimum descent or decision altitude provided in each approach chart.

Although some approaches had the height above the runway threshold published, the difference between the overall aerodrome elevation and published minimum descent altitude was used to align with the weather reports were assumed to use the aerodrome elevation as a datum.

The determination of cloud base from METARs and SPECIs is discussed in depth in the section 'Calculation of the weather ceiling' on page 43.

Wind vector calculations and comparisons

Tailwind and crosswind values were determined by calculating the resultant vector of the forecast wind conditions relative to the runway direction. This was in accordance with the following equations:

$$Tailwind = -\mathrm{V}_{WIND_{max}} \times cos(\Delta\theta), Crosswind = \mathrm{V}_{WIND_{max}} \times sin(\Delta\theta)$$

Where $\mathrm{V}_{\mathrm{WINDmax}}$ is the maximum forecast wind speed (the average is used if no gusts are forecast), and Δθ is the smallest angular difference between the wind direction and runway direction (re-aligned to true north as per the forecast[25]).

When wind direction was reported as variable, the tailwind and crosswind components were assumed to both be the worst-case maximum gust value forecast.

Aircraft crosswind and tailwind limits were adjusted if particular conditions were reported or approaches were required as follows.

- Aircraft tailwind limits
  - Adjusted to zero for CAT III b approaches.
- Aircraft crosswind limits
  - Unaltered for runway width of 45 metres and above
  - Runway width less than 45 metres used the following information if available
    - 30 metre wet runway crosswind limit – when heavy precipitation (rain, snow or drizzle) was reported of forecast
    - 30 metre dry runway crosswind limit – for dry runways less than 45 metres

Note that if all crosswind values were not available for an aircraft model or category, the next most stringent crosswind limits were applied. For example, if the wet 30 metre crosswind limit was not available for a wet runway under 45 metres wide, the 30 metre dry limit would be applied, followed by the maximum demonstrated crosswind if this was not available.

---

[25] All forecasts and observations have wind directions published relative to true north.

At the time of writing, no other limits to aircraft crosswind or tailwind were applied. However, the algorithms have been designed to allow modification of these effects if desired in the future.

Reported threshold runway wind shear was not taken into account at the time of writing, although reported thunderstorms at the aerodrome was considered to produce conditions below safe landing limits.

Comparison of landing minima criterion and selection of instrument approach procedures

Table 26 shows a sample of two SPECI type weather reports being compared to the landing minima applicable to instrument approaches available to a Boeing 737-800. This example is based on data from Adelaide and Melbourne Airports, some data has been altered in some cases for demonstration purposes.

All identified options from the data table ‘Aircraft approach minima’, as shown in Figure 10 were combined with each report. Data sourced from this table is contained under the left hand side columns labelled ‘Instrument approach and aircraft data’.

Data for comparison is shown in the right hand side of Table 26, with data from instrument approach procedures (IAP), aircraft specifications (AC) and METAR and SPECI (M) reports being listed. Modifications to this data are discussed in ‘Selection and calculation of landing minima criterion data fields’ from page 59.

The columns in Table 26 labelled ‘PD’ and ‘Weighting’ were calculated values used in the determination of conditions below the landing minima and the selection of a single instrument approach for each METAR and aircraft combination. These are discussed below.

Table 26: Sample selection of instrument approach for each METAR or SPECI

Instrument approach (IAP) and aircraft data				METAR or SPECI comparison with instrument approach and aircraft data																
				Visibility			Runway visual range			Ceiling			Crosswind			Tailwind				
Aircraft (AC)	**RWY**	**Approach**	**METAR phenomena**	**IAP**	**M**	**PD**	**IAP**	**M**	**PD**	**IAP**	**M**	**PD**	**AC**[26]	**M**	**PD**	**AC**	**M**	**PD**	**Below landing minima**	**Weighting**
Boeing B737-800	23	VOR/ DME	No thunderstorms, or heavy precipitation	4200	2500	-41	N/A	N/A	N/A	750	1100	47	33	3	91	10	-26	360	Yes	61
	23	RNAV RNP		1700		47				480		129		3	91		-26	360	No	231
	23	ILS I		800		213				250		340		3	91		-26	360	No	608
	05	RNAV RNP		2200		14				421		161		3	91		26	-160	Yes	178
	05	RNAV GNSS		2800		-11				500		120		3	91		26	-160	Yes	113
	34	RNAV RNP	No thunderstorms, or heavy precipitation	1900	100	-95	N/A	N/A	N/A	257	4998	1845	33	5	100	10	-5	150	Yes	1785
	27	ILS I		800		-88	550	363	-34	176		2740		0	85		-1	110	Yes	2735
	16	ILS I		800		-88	550	125	-77	206		2326		0	100		5	50	Yes	2274
	16	ILS II		N/A		N/A	300	125	-58	98		5000		0	100		5	50	Yes	4967
	16	ILS IIIa		N/A		N/A	175	125	-29	48		10312		0	100		5	50	Yes	10309
	16	ILS IIIb		N/A		N/A	75	125	67	0		49980		5	100	0	5	-500	Yes	50017
	09	RNAV RNP		2300		-96	N/A	N/A	N/A	397		1159		0	85	10	1	90	Yes	1091

This example is shown for two different locations, each with five instrument approach procedures (IAPs) applying to the selected aircraft. Data from each of the five rows would be applied to every METAR and SPECI in the same fashion. Note that even if other runways exist, the algorithm has

---

[26] The selection of aircraft crosswind limits was based on a runway width above or below 45 metres, and if any heavy precipitation (snow, rain or drizzle) was observed (or forecast).

excluded these based on none of the instrument approach procedures applying to the selected aircraft.

For example, consider the top five rows showing instrument approach options for runways 05 and 23 for a Boeing 737. Even if the airport had an additional runway 18/36, only runways 05 and 23 are combined for comparison because published instrument approach procedures applicable to the aircraft exist, and in this example do not exist for runway 18 and 36. These comparisons are covered in further depth in the section ‘Instrument approaches linked to aircraft models’ on page 48.

Calculation of percentage difference between METAR and SPECI report values and landing minima criteria

The columns labelled ‘PD’ are the calculated percentage difference between the published alternate minima threshold, defined by the following equations, where ‘Vis’, ‘Ceil’ and ‘RVR represent values of visibility, ceiling and RVR calculated separately, and ‘Wind’ represents values of either crosswind or tailwind, also calculated separately of one another:

$$PD_{Vis,Ceil,RVR} = \frac{F_{Vis,Ceil,RVR} - IAP_{Vis,Ceil,RVR}}{IAP_{Vis,Ceil,RVR}} \times 100, PD_{Wind} = \frac{AC_{Wind} - F_{Wind}}{AC_{Wind}} \times 100$$

The percentage difference was used for two purposes:

- showing if the minima is breached - a negative value indicates conditions below a minima criterion[27] (shown in red shading in Table 26)
- the difference between published thresholds and respective weather forecast or observation values - the magnitude of percentage difference. Variations in these values were used to select a single set minima values.

Assessment of METARs and SPECIs below the landing minima

Using the percentage difference values calculated for each available approach at an aerodrome, each approach was assessed for conditions above or below the landing minima. An observation-approach pair were recorded as being above the minima if all of the following criteria applied.

- PD Visibility > 0 or PD RVR > 0, and
- PD Ceiling > 0 and
- PD Crosswind > 0 and
- PD Tailwind > 0 and
- no thunderstorms

For example, Table 26 shows two observation-approach pairs of conditions above the landing minima for the ILS and RNAV RNP approaches for runway 23 as all percentage differences, and therefore individual minima criterion were above the landing minima thresholds.

Worst case conditions assumed

When fields were missing, preventing an assessment of the conditions, these records were assumed to produce a safe outcome with a blank assessment recorded.

---

[27] The two equations for percentage difference were selected to always have a negative value when conditions were below the minima. This was required because the instrument approach procedures are based on conditions above a minimum value, and the aircraft crosswind and tailwind limitations being a maximum value.

Instrument approach weighting and selection

In the majority of cases when conditions were above the landing minima, flying more than one approach was possible. However, it was desired that only one approach be selected to show the potential dynamics of runway and approach use. Additionally, when conditions were observed below the landing minima for every approach, an approach was still nominated to assess the severity of the conditions relative to the minima breached.

A single weighted value for every forecast and alternate minima combination was calculated according to the following equation:

$$Weighting = \max(PD_{RVR}, PD_{Vis}) + PD_{Ceil} + 0.1 \times (PD_{Crosswind} + PD_{Tailwind})$$

This weighting was nominally chosen, with different weightings able to be added to each individual percentage difference if desired, allowing refinement of this model. A larger weighting generally indicates that the instrument approach was more accurate, with a greater difference between reported observations and landing minima criterion.

The largest of either RVR or visibility percentage differences were used in the calculation to provide a single weight for visibility fields, and due to only one of these values being required for the landing minima assessment.

This weighting gives priority to variations in the published visibility and ceiling minima over wind conditions approaching aircraft limits, which were reduced to 10 percent of the original values. The column labelled ‘Weighting’ in Table 26 demonstrates some examples of these calculations showing the overall value for each landing minima set.

Note that thunderstorms were not included in the weighting due to these having no relative difference between different alternate minima options.

Selection of instrument approaches for METAR or SPECI to simulate efficient operations

To select only one approach for every METAR or SPECI, a ranking was calculated using the assessment below landing minima, calculated weighting for each report-approach pair and individual percentage differences. This was used to rank the relative ‘suitability’ of each landing minima against each of the observed conditions. The final output of this process was an assessment for every METAR and SPECI linking to a specific instrument approach procedure.

The selection of a single approach was performed to simulate the effect of decision making by operational personnel. However, the framework of this algorithm allowed for the implementation of other decision-making hierarchy for use in future analyses.

The selection of an instrument approach to allow the most efficient operations was assumed to provide the most realistic scenario with regard to reflecting actual real-world operations. This was applied when conditions were above at least one approach landing minima. This provides the smallest margin between the published landing minima and conditions, continually increasing the precision of the selected approach with deteriorating conditions, down to a point where no approach was possible above the landing minima. The objective of this was to simulate the highest approaches taken for the highest possible acceptance rate. This also allowed differentiation between the severity of different observations’ situations.

When conditions were below the landing minima for every approach, an approach as close as possible to the conditions was selected.

The following ranking hierarchies were used to select an approach simulating efficient operations and the most accurate approach when observations below the landing minima for all approaches were reported.

Above landing minima selection: at least one suitable approach above landing minima – least stringent conditions selected

1. Select approaches with conditions above landing minima, then
2. Select approach with *lowest* selection weighting, then
3. Select *lowest* percentage differences in the following order: (PD Vis, PD RVR,PD Ceiling, PD Crosswind, PD Tailwind, thunderstorms)

The top five rows of data on Table 26 show an example of conditions observed above the landing minima for at least one suitable instrument approach procedure. Using the hierarchy above, the runway 23 RNAV RNP approach has been selected (highlighted in blue). This is because all minima criterion were above the landing minima, narrowing the selection to between the runway 23 ILS category I approach and RNAV RNP. The RNP approach was selected because it had the least lowest selection weighting, which was assumed to be the most likely approach selected by flight crew.

Below landing minima selection: No suitable approach for landing – closest conditions selected

1. Select the approaches with the least number of landing minima conditions breached, then
2. Select the approach with the *highest* selection weighting
3. Select the *highest* percentage differences in the following order: (PD Vis, PD RVR, PD Ceiling, PD Crosswind, PD Tailwind, thunderstorms)

For conditions where no landing above any landing minima was possible, the approach providing the smallest margin between observed conditions and the landing minima thresholds was selected. The lower portion of Table 26 shows an example where all approaches had minimum criterion below the landing minima. Using the hierarchy for all conditions below the landing minima, the approaches with only one landing minima breached were selected, narrowing the selection to five approaches. This was followed by the selection of the approach with the highest weighting, with runway 16 ILS Category III b approach being selected (highlighted in blue).

At the time of writing, it was considered that these selection weightings and hierarchies required refinement and comparison with actual operations to create a more accurate model. In addition, there may be other operational aspects such as accounting for the protection of ILS zones that could be considered in future analysis. Another decision making criteria could be a visual preference approach, which may allow for the analysis of the least stringent approach required for landing and alternate minima.

However, the overall assessment of reported conditions being above or below the landing minima was unaffected by these hierarchies.

Assessment of TAFs and TTFs against alternate minima criteria

This section describes the pairing process to link aircraft specific alternate minima (described on page 51) to each TAF and TTF sub-segment (described from page 52 and page 58). For every TAF and TTF sub-segment, one set of alternate minima criteria were selected for each nominated aircraft. Although examples presented are for TAFs, the process followed was the same for TTFs. At the time of writing, the alternate minima criteria for TAFs and TTFs were based on AIP ENR 58, which is summarised in ‘

Appendix C – Assessment of a safe forecast' on page 107.

To select the alternate minima criteria, a number of comparisons were performed between each forecast sub-segment, and the list of possible runway, aircraft and alternate minima combinations.

Table 27 shows a sample of two forecast sub-segments compared to the alternate minima applicable to a Boeing 737-800 at the forecast aerodrome. This follows a similar method described to select the landing minima criteria (described on page 58). This table uses the same data from Table 20, Table 22 and Table 23 to aid understanding. Although this example is based on Adelaide Airport, data has been altered for demonstration purposes.

On the left side of the table, aircraft and instrument approach data labels are shown indicating the aircraft model, runway and aircraft category, where 'S' indicates the special alternate minima. All identified options from the data table 'Aircraft alternate minima runway', as shown in Figure 10 were combined with each forecast sub-segment.

In this example, three different alternate minima options for two different runways apply to the selected aircraft, and the data from each of the three rows was applied to every sub-segment in the same fashion. In the same method described for matching instrument approach procedures to METARs and SPECIs, even if other runways exist, the algorithm has excluded these based on none of the instrument approach procedures applying to the selected aircraft.

A discussion for the selection of these rows is contained in the section 'Determining possible alternate minima applicable to an aircraft' on page 51.

Table 27: Sample selection of alternate minima values for each TAF and TTF sub-segment

Instrument approach (IAP) and aircraft data			Fore-cast data (F)		Forecast comparison with instrument approach and aircraft data													
					Visibility			Ceiling			Crosswind			Tailwind				
Aircraft (AC)	**RWY**	**Category**	**Sub-segment**	**Thunderstorms**	**IAP**	**F**	**PD**	**IAP**	**F**	**PD**	**AC**[28]	**F**	**PD**	**AC**	**F**	**PD**	**Below alternate**	**Weighting**
Boeing B737-800	23	D	TAF- FM 1	No	7000	9000	22	1313	1200	-9	33	0	100	10	-17	270	Yes	50
	23	S			4000		125	850		41		0	100		-17	270	No	203
	05	D			7000		22	1313		-9		0	100		17	-70	Yes	-47
	23	D	PROB 2	No	7000	500	-93	1313	4998	280	33	0	100	10	-17	270	Yes	224
	23	S			4000		-88	850		488		0	100		-17	270	Yes	437
	05	D			7000		-93	1313		280		0	100		17	-70	Yes	127

The column in Table 27 labelled ‘Sub-segment’ was included to provide a reference to the examples as mentioned. ‘PROB 2’ is also referred to as ‘P 2’ in Figure 15 below. The columns to the right of the ‘Forecast data’ show comparisons between published alternate minima thresholds. This was sourced from instrument approaches (labelled IAP) and aircraft (labelled AC) and forecast sub-segments (labelled F).

The forecast presence of thunderstorms, combined with the percentage differences from each of the other alternate minima criteria were used to determine if conditions were forecast below any of the alternate minima thresholds. This is illustrated in the ‘Below alternate’ column of Table 27.

For permanent sub-segments, if any of the values of visibility, ceiling or wind were not present, an assessment of the alternate minima was not performed. For temporary and probabilistic sub-segments, at least one of these values was required for an assessment. Due to the formation process for these sub-segments, this was often dependent on the coinciding permanent sub-segment, as described in the section ‘TAF temporary and probabilistic change segment adjustments’ on page 56. The result of these processes is that the forecast sub-segments are excluded from the analysis.

Selection of alternate minima using weighted values

As it was possible for multiple alternate minima value sets to apply to a given forecast, a rank was calculated for every alternate minima value set for every forecast sub-segment. This used a decision-making hierarchy to establish the relative ‘suitability’ of each alternate minima against each of the forecast conditions.

---

[28] The selection of aircraft crosswind limits was based on a runway width above or below 45 metres, and if any heavy precipitation (Snow, rain or drizzle) was forecast.

This was performed to simulate the effect of decision making by operational personnel, and also allowed the comparison of only one set of values for further analysis. However, the framework of this algorithm allowed for the implementation of other decision-making hierarchy to be implemented for use in future analyses.

A single weighted value for every forecast and alternate minima combination was calculated according to the following equation:

$$Weighting = PD_{Vis} + PD_{Ceil} + 0.1 \times (PD_{Crosswind} + PD_{Tailwind})$$

This weighting was chosen to give priority to variations in the published alternate visibility and ceiling minima (such as the special alternate minima) over wind conditions approaching aircraft limits. The column labelled ‘Weighting’ in Table 27 demonstrates some examples of these calculations showing the overall value for each alternate minima set.

Note that forecast thunderstorms were not included in the weighting due to these having no relative difference between different alternate minima options.

The first priority for alternate minima selection was conditions above the alternate minima, meaning that if a crosswind or tailwind exceeding the published aircraft limits existed, alternate minima for a runway with a more favourable wind direction will be selected.

If only one alternate minima set existed above the alternate minima, this was selected. Table 27 sub-segment labelled ‘TAF-FM1’ shows this scenario, with the weighting and alternate minima columns highlighted in ‘blue’.

The alternate minima set with the largest weighting were selected in cases where all sets were below the alternate minima, or where multiple options existed for forecast conditions above the alternate minima. Table 27 sub-segment labelled ‘PROB 2’ shows one of these scenarios (selection again highlighted in ‘blue’), with all options being below the alternate minima.

Choosing the largest weighting as calculated above has the effect of selecting the lowest possible alternate minima for visibility and ceiling. For example, if special alternate minima apply to an instrument approach, and the current forecast would allow a landing on this runway above that minima, this was selected over the standard alternate minima applicable to that category of aircraft.

Comparison of overlapping TAF and TTF forecasts with METAR and SPECI reports

To assess the reliability of each forecast sub-segment, METAR and SPECI reports representing observations within the forecast sub-segment start and end times were synchronised. These comparisons use alternate minima assessments performed on every TAF and TTF (discussed on page 67), and the landing minima assessments performed on every METAR and SPECI (discussed on page 58).

At the end of this process, comparisons between all combinations of time overlapped forecasts and reports were performed. The comparisons stored in this table were used to assess the predictions of each forecast sub-segment in *isolation*. This data was subsequently used to assess the overall reliability of the TAF or TTF, and is discussed in depth in the section ‘Overall forecast assessments for operational reliability’ on page 77.

It is important to note that at this stage, the interactions of temporary forecast sub-segments with other time overlapping temporary and permanent sub-segments were not assessed, and therefore do not necessarily infer the overall operational impact of the forecast.

Release times of METAR and SPECI reports used to approximate changes in observed conditions

To combine and compare METAR and SPECI data with forecast information, nominal start and end times were assigned to each METAR and SPECI report as follows:

- Start time: The METAR or SPECI release time was used to estimate the commencement of conditions reported in that report.
- End time: The release of the next METAR or SPECI report was used to estimate the end of reported conditions, and the start of the next conditions.

If there was more than 60 minutes in between the release of consecutive weather reports, the earliest report was to be flagged in a separate table for exclusion from analysis, as described in 'Time exceeding typical duration between forecasts and reports' on page 83.

At the time of writing, there were *no adjustments of time performed* on METAR and SPECI reports, such as accounting for averaging processes of wind, visibility and cloud base, or accounting for assurance related release time delay when improvements in weather conditions are observed. It was expected that error between actual conditions and those approximated used in this study would exist associated with always using the release times to represent the timing of all conditions. In particular, error was expected in the timing of the weather parameters noted above.

However, at times of particular interest in this study, where weather is poor or transient, the frequency of reports (generally in the form of SPECIs) is expected to be higher, thereby increasing the general confidence in the timing of reported values in degrading conditions, with an expected 10 minute lag when conditions are improving. Furthermore, due to the study focusing on whether conditions are forecast or observed below the alternate and landing minima, the precise nature of the exact conditions are less critical to support the overall analysis, even though it is expected that discrepancies will be present when comparing actual values.

Combining TAF and TTF forecasts with METAR and SPECI reports

Comparing forecast data with weather report information required these data sources to be combined. This process was similar to that performed for creating temporary TAF and TTF sub-segments, as described on page 52, with regard to the selection of start and end times.

A comparison was performed between every forecast sub-segment and overlapping METAR and SPECI reports (with start and end times as discussed). Figure 15 illustrates how the alignment of times between forecast sub-segments and weather reports was achieved by creating timing divisions (shown as vertical blue dashed lines) whenever a change in forecast sub-segment or weather reported occurred. For example, the permanent forecast sub-segment labelled FM 1 (in yellow) spans over the two weather reports labelled M5 and M6 (shown in pink). The formed forecast sub-segment – report comparison pairs created are labelled C3 and C4 and are shown in light blue. Another example shows the METAR report labelled M8 (shown in pink) being used to compare with two forecast sub-segments labelled FM 1 – FM 2 (in red), and 'FM 1 – FM 2 – BECMG' (in purple). In this case, the timing point was created in the middle of report M8 to create comparison pairs C6 and C7.

Note that the permanent and overlapping temporary and probabilistic forecast sub-segments remain separate in this process, with independent comparisons for every temporary, probabilistic and permanent sub-segment.

Table 28 shows an example synchronisation of times for a number of forecast sub-segment – report pairs shown in Figure 15. The forecast sub-segment data is also a continuation of the example in Figure 12 and Table 21. Note also that the synchronisation process was very similar using permanent sub-segments to dividing temporary TAF segments into sub-segments.

Figure 15: Synchronisation of TAF and TTF forecasts with METAR and SPECI reports (sample data continued from Figure 12)

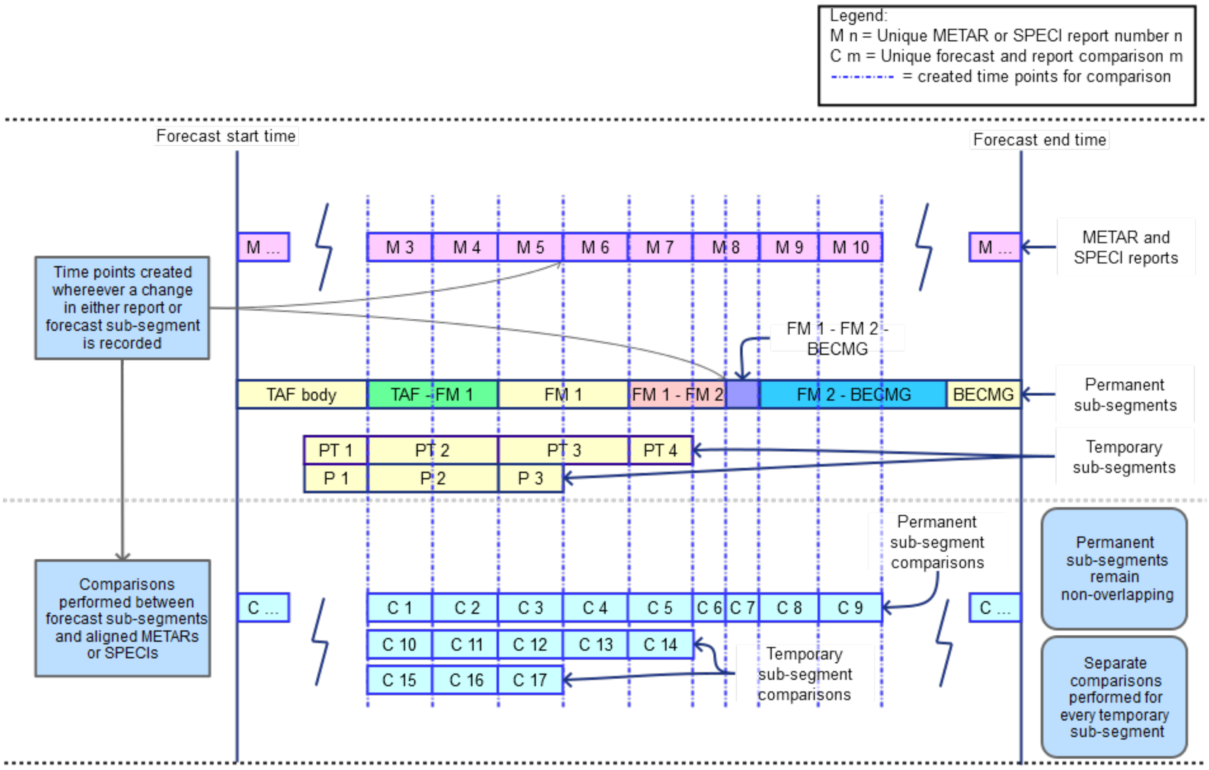

Table 28: Sample synchronisation between weather forecast (TAF and TTF) sub-segments and weather reports (METAR and SPECI) as shown in Figure 15

METAR or SPECI release times (new)		Forecast sub-segment times			Comparison times		
Label	**Release time**	**Label (from Table 21)**	**Start time**	**End Time**	**Label**	**Start time**	**End Time**
Permanent sub-segments							
…	…	…	…	…	…	…	…
M 5	14:30	FM 1	14:30	15:30	C 3	14:30	15:00
M 6	15:00				C 4	15:00	15:30
M 7	15:30	FM 1 – FM 2	15:30	16:15	C 5	15:30	16:00
M 8	16:00				C 6	16:00	16:15
		FM 1 – FM 2 – BECMG (F1- F2- B)	16:15	16:30	C 7	16:15	16:30
M 9	16:30	FM 2 – BECMG	16:30	19:00	C 8	16:30	17:00
M 10	17:00				C 9	17:00	…
Temporary sub-segments (PROB TEMPO)							
M 5	14:30	PT 3	14:30	15:30	C 12	14:30	15:00
M 6	15:00				C 13	15:00	15:30
M 7	15:30	PT 4	15:30	16:00	C 14	15:30	16:00
M 8	16:00						
Temporary sub-segments (PROB)							
M 5	14:30	P 3	14:30	15:00	C 17	14:30	15:00
M 6	15:00						

Comparisons between forecasts and reports

For every forecast sub-segment and report pair, comparisons between each of the weather forecast parameters were performed in addition to the comparison of operational minima. Two types of comparisons were performed based the type of data – those for continuous variables, and comparisons based on binary fields, such as the existence of a weather phenomenon.

The data table created in this process provides the formal link between all forecasts and reported observations. This enabled the overall assessment of each forecast (discussed on page 77), and also contains the record of every forecast sub-segment compared to synchronised METARs and SPECIs, regardless of whether these were excluded due to having no operational effect in the overall assessment process. This table can be reconfigured to assess every element individually without contingency time overlaps, which may be useful for analysis of different forecast segment types.

Table 29 shows an example of both of these types of comparisons between forecast sub-segments and weather reports. This table is also a continuation of data linked to Figure 15, and forecast data (as per the previous examples) to Figure 12. Columns labelled ‘F’, ‘M’ and ‘C’ represent forecast (TAF or TTF), report (METAR or SPECI) and comparison data respectively, where colour coding in the left three columns corresponds to elements in Figure 15 and Table 28.

Table 29: Sample comparisons between forecasts sub-segments and reported observations

Figure 15 reference			Comparisons between forecasts and observations																		Operational assessment		
Forecast	METAR / SPECI	Comparison	Visibility			Ceiling			Wind angle			Wind speed			Thunderstorms			Fog			Below alternate	Below landing	Rating (Figure 16)
F	M	C	F	M	C	F	M	C	F	M	C	F	M	C	F	M	C	F	M	C	F	M	C
FM 1	M 5	C 3	9000	2500	-6500	1200	1100	-100	260	230	30	14	26	-12	No	No	4	No	No	4	No	No	4 (CR)
PT 3		C 12	2000		500	1000		100	VRB		N/A	20		-6	Yes		3	No		4	Yes		3 (FA)
P 3		C 17	500		2000	4998		-3898	220		10	17		-9	No		4	Yes		3	Yes		3 (FA)
FM 1	M 6	C 4	9000	2000	-7000	1200	1000	-200	260	200	60	14	26	-12	No	Yes	2	No	No	4	No	Yes	2 (MISS)
PT 3		C 13	2000		0	1000		0	VRB		N/A	20		-6	Yes		1	No		4	Yes		1 (HIT)
FM 1 – FM 2	M 7	C 5	9000	2000	-7000	1200	1100	-100	260	250	10	14	35	-21	No	Yes	2	No	No	4	No	Yes	2 (MISS)
PT 4		C 14	2000		0	1000		100	VRB		N/A	20		-15	Yes		1	No		4	Yes		1 (HIT)
FM 1 – FM 2	M 8	C 6	9000	1000	-8000	1200	500	-700	260	230	30	14	5	9	No	No	4	No	Yes	2	No	Yes	2 (MISS)
F1 – F2 – B	M8	C 7	1500	1000	-500	1200	500	-700	210	230	50	16	5	11	No	No	4	Yes	Yes	1	Yes	Yes	1 (HIT)

The middle section of Table 29 labelled ‘Comparisons between forecasts and observations’ shows typical comparisons performed between the reported and forecast parameters that are discussed in the following sections. Note that these columns were produced to enable further analysis and assessment, and were not used for operational assessment.

The columns ‘Below alternate’ and ‘Below landing’ under the label ‘Operational assessment’ refer to the already assessed minima conditions for forecasts and reports as discussed in the sections ‘Assessment of TAFs and TTFs against alternate minima criteria’ on page 67, and ‘Assessment of METARs and SPECIs below the landing minima’ on page 65. Samples of these calculations, including links to aircraft models and instrument approaches are contained in Table 26 and Table 27, with selected outcome (as would be inserted in the respective Table 29 columns) shown in ‘Blue’.

Data under the column labelled ‘Rating (Figure 16)’ contains a categorised outcome of the operational assessment of the sub-segment compared with each report. This is discussed in depth in the section ‘Comparing binary forecast-report elements related to operational impact’ on page 76.

The horizontal divisions separated by the double line in Table 29 denote comparison groups with the same start and end times, as can be observed in Figure 15. Furthermore, the comparison of all operational assessments within a group forms the basis for the assessment of the forecast reliability, discussed in the section ‘Overall forecast assessments for operational reliability’ on page 77.

Comparison of quasi-continuous values

Comparisons between forecasts and synchronised reports for quasi-continuous[29] values wind, visibility and cloud base were performed by recording the difference between forecast value and reported value. This is shown in the comparison columns (labelled ‘C’) in Table 29.

In the same way that the percentage difference calculations were conducted for the calculation of these parameters for alternate and landing minima, the calculations were performed to be negative when the forecast condition was less stringent than the reported observation. For example, the difference in wind speed would be reported as negative if the forecast wind speed was *less* than the reported wind speed, and conversely, the difference in visibility would be negative if the forecast stated figures *more* than the reported visibility. However, note that these negative values do not necessarily indicate conditions below the minima.

Binary value comparisons

Binary fields were used to store two main forms of information for the analysis of weather reliability:

1. Weather phenomena predictions
2. The operational impact of forecasts and observations, for example, reported conditions above or below the landing minima.

To compare one binary value from a forecast (TAF or TTF) and one binary value from a synchronised report (METAR or SPECI), four categories applied, corresponding to one of four possible permutations. Figure 16 shows a binary comparison matrix, which defines the categories used for every binary forecast – report element pair.

---

[29] These values were assumed to be quasi-continuous due to the fidelity of this information being dependent in some cases on the magnitude of the values. For example, visibility reported below 1000 will be given to the nearest X feet, compared with being reported above 5000 metres, where it will be given to the nearest Y feet.

Figure 16: Forecast predictions compared to reported observations

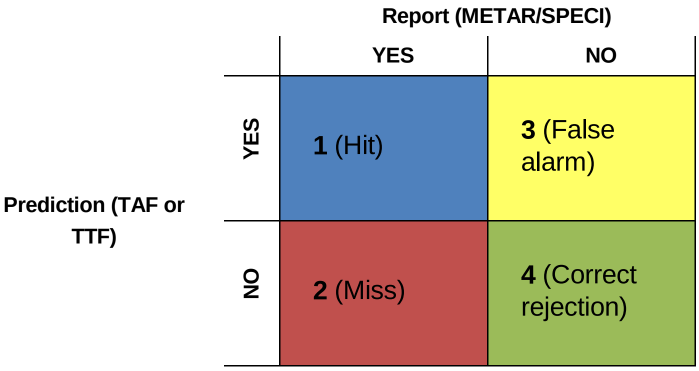

		Report (METAR/SPECI)	
		YES	NO
Prediction (TAF or TTF)	YES	**1** (Hit)	**3** (False alarm)
	NO	**2** (Miss)	**4** (Correct rejection)

Weather phenomena predictions

The presence of any weather phenomena and descriptors were compared between each forecast-report pair. A category was recorded in accordance with Figure 16.

For example, if thunderstorms were reported but were not forecast, the number ‘2’ would be inserted under thunderstorms for the forecast-report pair. Sample results of these outcomes are shown under the comparison column ‘C’ for thunderstorms and fog in Table 29.

Comparing binary forecast-report elements related to operational impact

Binary field comparisons were performed between the calculated landing and alternate minima states (described on pages 65 and 69 respectively) of reports and forecast sub-segments respectively. The alternate minima state was also derived for each report (METAR or SPECI) segment using the merged forecast alternate minima values. The same techniques as that for forecasts were used to determine if reported conditions were below the *alternate* minima.

When combined with all forecast sub-segment – report pairs from a single forecast, these operational comparisons were used as the basis for assessments of the overall forecast reliability, and is described on page 77.

At the time of writing, two binary forecast-report comparisons relating to operational impact were performed:

- *Alternate* minima state of TAF and TTF sub-segments compared with *landing* minima state of METARs and SPECIs
- *Alternate* minima state of TAF and TTF sub-segments compared with *alternate* minima state of METARs and SPECIs

At the time of writing, an additional comparison was considered, but not performed:

- *Landing* minima state of TAF and TTF sub-segments compared with *landing minima* state of METARs and SPECIs

The comparison of landing minima for both forecasts and reports was considered to allow additional fidelity in understanding the use of the landing minima as a threshold for TAFs and TTFs instead of the alternate minima. This may be useful for future analyses.

Comparison of forecast alternate minima and report landing minima states

Forecast conditions below the alternate minima were assumed to impose an operational requirement, and observed conditions below the landing minima was assumed to be the limit of an aircraft to conduct a safe landing[30].

---

[30] However, given certain conditions, it is expected that it is probably favourable to land in conditions that are slightly below the published landing minima thresholds – where it would pose a

The comparison of these two states was performed to indicate the potential operational effect of planning based on the alternate minima state of the forecast, and the resultant reported conditions affecting the ability to land. This measure is intended to provide an indication of the effectiveness of the forecasting safety system by taking account of the buffer between the alternate and landing minima.

For example, if conditions were forecast below the alternate minima, it is expected that a contingency plan would be able to be implemented by a flight crew, such as taking on additional holding fuel or planning for an alternate aerodrome. If on arrival, conditions were reported below the landing minima, this comparison would be recorded as a ‘Hit’ according to Figure 16. Note that if conditions were reported above the landing minima (even if below the alternate minima), the forecast sub-segment – report pair would be recorded as a ‘False alarm’ for this comparison.

In contrast, if forecast conditions were above the alternate minima and reported conditions were below the landing minima, this was recorded as a ‘Miss’. *Forecast* conditions above the *alternate* minima and *reported* conditions above the *landing* minima were recorded as a ‘Correct rejection’, even for *reported* conditions between the landing and alternate minima states.

An example of these resultant comparisons is shown in the column labelled ‘Rating (Figure 16)’ of Table 29. These are colour coded to match Figure 16.

Comparison of forecast alternate minima and report alternate minima

The comparison of the alternate minima state between every forecasts sub-segment and report was performed to allow the assessment of states *between reported* landing and alternate minima. The same comparison method was used as for alternate and landing minima comparisons described above. This used the calculated *alternate* minima state of the METAR or SPECI from the same alternate minima thresholds used in the assessment of the paired TAF or TTF sub-segment.

For example, if the forecast was calculated to be above the alternate minima, however was *reported* to be below that same threshold, this would be recorded as a ‘Miss’ (again referring to Figure 16). Values for these comparisons were inserted into the same data table as the comparison of alternate and landing minima.

The assessment of states reported *between* landing and alternate minima were occurred as part of the overall forecast assessments, and are discussed on page 77.

Overall forecast assessments for operational reliability

This section describes the process of assessing the overall operational effect of a TAF or TTF when compared to METARs and SPECIs that are synchronised for each time of prediction. At the end of this process, all overlapping permanent and temporary sub-segments (as described from pages 52 and 58) are resolved to show the operational reliability over the duration of the forecast validity period.

The main operational reliability indicator used was the forecast alternate state compared to the landing minima state of the reported observation on page 76, taking into account all overlapping segments within the same TAF or TTF.

The assessment of operational reliability for every individual TAF or TTF involved two key processes described under the following sub-headings:

- Formation of non-overlapping reliability timeline for each TAF and TTF
- Summary of operational reliability of each TAF and TTF

---

greater perceived risk to continue flying rather than land. This may warrant further research to understand any practical breaches in landing minima that occur.

Formation of non-overlapping reliability timeline for each TAF and TTF

This process aims to simulate the operational effect of a forecast throughout the validity period. To achieve this, all permanent and overlapping temporary segments were evaluated together to create a resolved single timeline by operational state (as shown in Figure 16 and discussed in the section 'Comparison of forecast alternate minima and report landing minima states' on page 76). This was important within each forecast, as it allowed the analysis of time when each operational state was present in every TAF and TTF.

To describe the reliability outcome timeline, a hypothetical scenario is presented. If a flight crew were planning to arrive at a given time, they would consider both the permanent and overlapping temporary and probabilistic segments active during the expected time of arrival. The assessment of whether a contingency plan is required for that arrival time (taking into account all time adjustments as discussed on pages 52 and 58) relies on whether *any* segment predicts conditions below the alternate minima, or more generally, triggers an operational requirement. When compared with actual conditions at the time of arrival, the reliability outcome is recorded as a 'Hit' or 'False alarm' whenever conditions are triggered or primed, and a 'Miss' or 'Correct rejection' when no forecast sub-segments predict conditions below the alternate minima (as shown in Figure 16).

The process of overall forecast assessment is similar to the flight crew planning process described above, except that instead of having a single arrival time for evaluation, an arrival *timeline* is produced showing the 'operational reliability' state throughout the entire TAF or TTF validity period. This creates the effect of recording the outcome of forecast predictions for each moment in time throughout each forecast validity period.

This process considers all forecasts in isolation, however, it forms the basis of the simulation of effect of retrieving one or more forecasts for the purposes of planning covering the same validity period. An in-depth discussion of this is contained in the section 'Methodology of results and analysis - simulation of the operational effect of forecast retrievals' on page 84.

Figure 17 provides an illustration of the processes used to resolve forecast comparison segments (as described on page 73), into a single timeline labelled 'resolved forecast comparisons' in the figure, and as described in the paragraphs above.

Figure 17: Resolved operational assessment by time for forecasts compared with reported observations – central example (red-dashed outline) continued from Figure 15

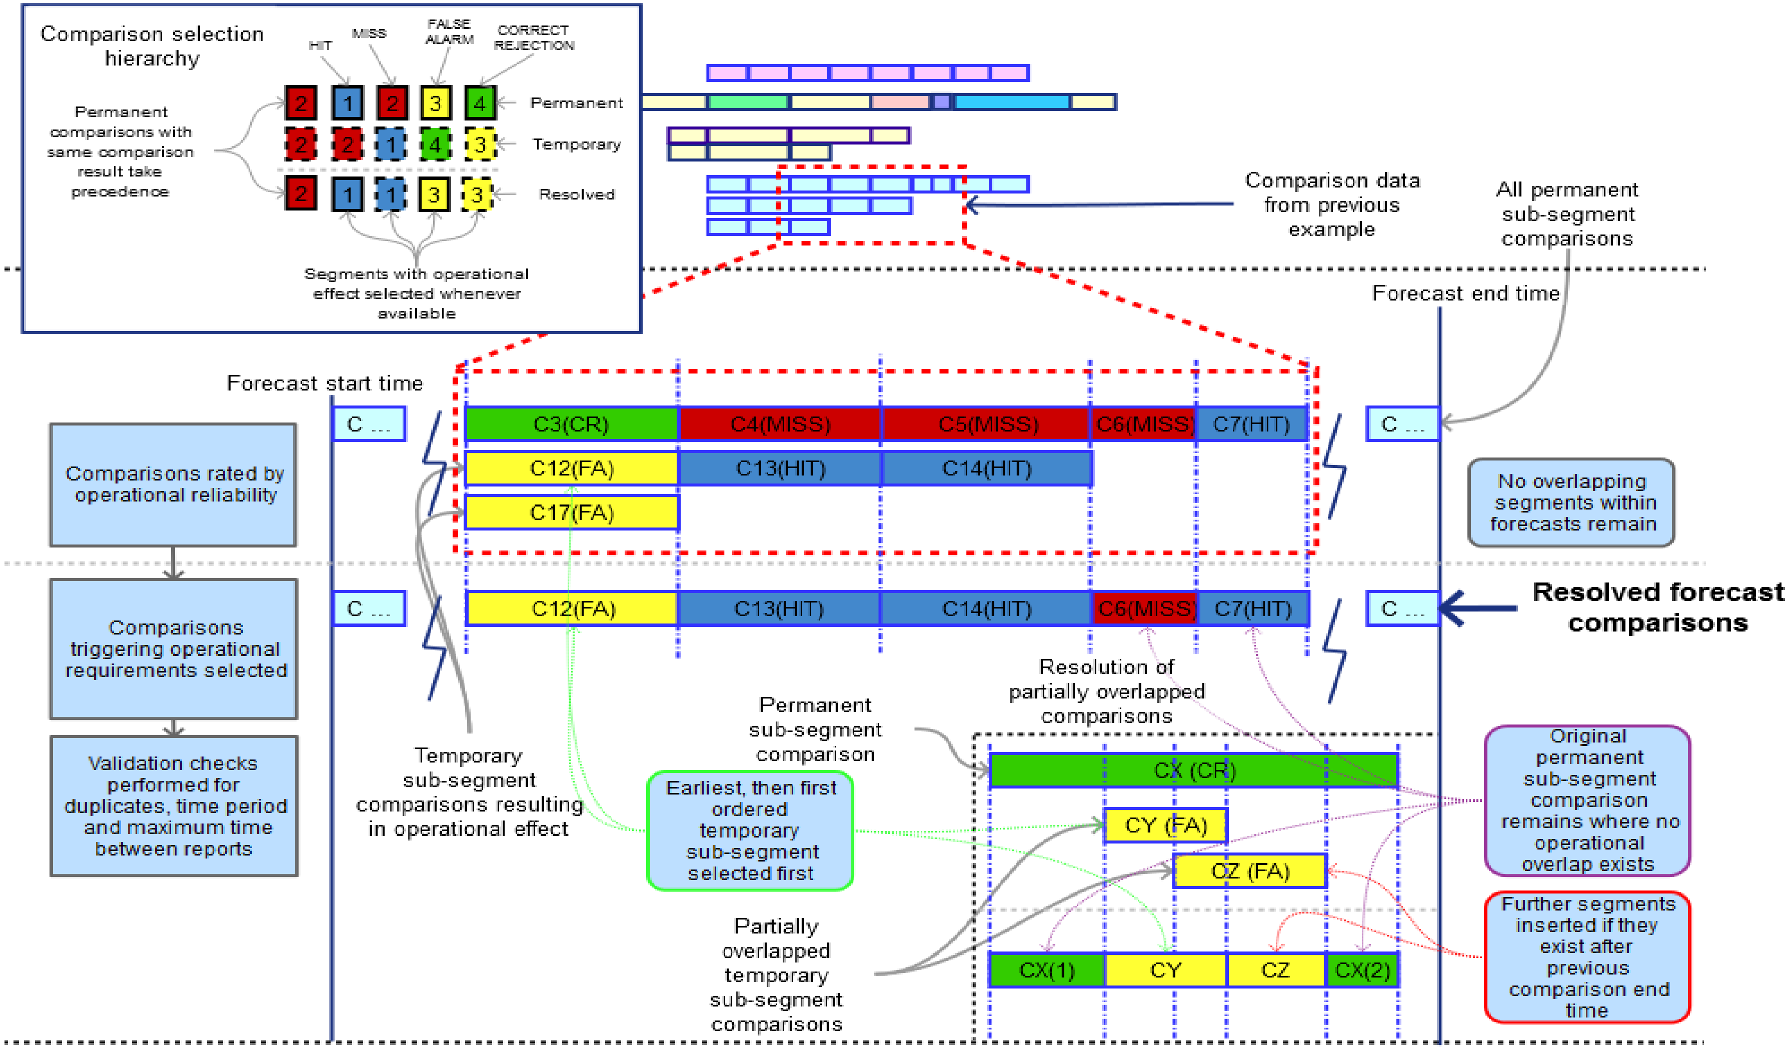

The main example depicted in the central area of Figure 17 is a continuation from Figure 15, outlined by red dashed lines. The labels on these data elements indicate the corresponding comparisons used. Colour coding and labels in brackets match the ‘Rating (Figure 16)’ column in Table 29.

Selection of comparison segments for forecast timeline of operational effect

To form a single operational timeline, individual comparisons were selected and ‘inserted’ into the timeline for every period within the forecast validity. When both permanent and temporary (including probabilistic) sub-segment comparisons co-exist, a selection hierarchy was used to determine the comparison sub-segment to be inserted in the timeline.

This is equivalent to a flight crew selecting one forecast segment to base decision making, in this case giving priority to any forecast segment creating an operational effect – and specifically creating alternate minima planning requirements.

The hierarchy used for selection of co-existing comparison segments is shown in the top left corner of Figure 17. This shows ‘hits’ always selected over ‘misses’ and ‘false alarms’ over ‘correct rejections’[31] For example, Figure 17 shows the co-existing permanent comparison labelled ‘C4 (MISS)’, and temporary comparison labelled ‘C13 (HIT)’. Comparison C13 was selected due to it creating an operational effect. Note that this simulates the selection of the ‘PROB TEMPO’ segment of the TAF, labelled ‘PT 3’, as shown in Figure 15 and Table 29.

When permanent comparison segments had no overlapping temporary comparison segments, this was inserted directly into the timeline, as shown in Figure 17 segments ‘C6’ and ‘C7’.

Where the operational outcome of all comparison segments was the same, such as all being recorded as a ‘Hit’, or ‘Miss’ (as shown), preference is given to the permanent sub-segment. When more than one temporary sub-segments co-exist creating an operational requirement, and the permanent segment does not, the first segment to be evaluated was arbitrarily[32] selected. This is shown by the selection of the yellow ‘False alarm’ comparison ‘C12’ in Figure 17.

Table 30 shows an example indicative of the data produced in the comparison resolution process illustrated in Figure 17, using the same data as previous examples. The forecast validity period and first and last rows of the comparison are new data in this example. This shows a single operational timeline for a 12-hour forecast. The comparisons labelled ‘C…’ are intended to indicate multiple comparison segments, assessed as a correct rejection, with no weather forecast below the alternate or observed below the landing minima.

The operational timeline data shows both the operational effect of forecasting and the reported observations. For example, it can be seen that operational requirements were triggered (as a false alarm) 30 minutes prior to conditions falling below the landing minima for a total of 90 minutes (shown by segments C13,14,6 and 7). This was mostly forecast at least below the alternate minima. However, during this time, there was a 15 minute envelope in which the forecast did not predict conditions below the alternate minima (segment C6), and was recorded as a ‘miss’.

---

[31] Note that hits and misses are mutually exclusive from false alarms and correct rejections.

[32] At the time of writing, the severity of the prediction of each forecast segment was not taken into account, as the focus of this analysis was on whether conditions were forecast below the alternate minima. A future iteration of this procedure may take into account the predicted severity as a consideration for the selection of comparison segments.

Table 30: Sample of single forecast assessment comparison data, based on expansion of selected resolved comparisons from Figure 17, Table 28 and Table 29

Forecast validity			Resolved forecast comparisons				
Start time	**End time**	**Duration**	**Comparison**	**Assessment**	**Start time**	**End time**	**Duration**
12:00	00:00	720	C…	Correct rejection	12:00	14:30	150
			C12	False alarm	14:30	15:00	30
			C13	Hit	15:00	15:30	30
			C14	Hit	15:30	16:00	30
			C6	Miss	16:00	16:15	15
			C7	Hit	16:15	16:30	15
			C…	Correct rejection	16:30	00:00	450

This information also provides further context in explaining the overall forecasting characteristics. In this example, it can be seen that adverse weather conditions *were* predicted with small holes in the forecasting, rather than being completely missed. Additionally, although a false alarm[33] is recorded, adverse weather was observed a short time later as opposed to no adverse weather being observed.

Resolution and selection of partially overlapping segments

In some cases, temporary forecast sub-segment comparisons did not completely align with permanent segments due to the sub-segments starting and finishing at different times. To account for this, new timing divisions were created to perform the operational assessment.

The lower right corner of Figure 17 shows an example of this process. The algorithm selects the earliest (or first in order if more than one start at the same time) overlapping temporary sub-segment comparison and inserts this into the timeline (labelled ‘CY’ in Figure 17). If more temporary sub-segment comparisons remain, the same process was followed from the end of the prior comparison segment, which was used as the start of the overlapping segment (‘CZ’). Any part of the permanent comparison segment (‘CX’) not overlapped by temporary comparisons was inserted into the timeline in these gaps (‘CX(1)’ and ‘CX(2)’ in the example).

Validation and summary of operational reliability of each TAF and TTF assessment

Following the process of creating a single timeline for each forecast, the timing data was summarised and checked for inconsistencies. The validation focused on forecast time duration as this was a critical component of this analysis, being the numerator for the majority of results and conclusions in analyses.

These validations were performed during the overall forecast assessment because this was the final ‘forecast-centric’ evaluation process, allowing a review of the evaluated data prior to it being used in a holistic simulation (described from page 84).

The validation of each resolved forecast-report comparison involved three types of assessments:

- Check-sums comparing the original decoded forecast validity period (discussed on page 33) with the resolved comparison segments (discussed on page 78)
- The assessment of duplicate or simulataneously released reports
- Identification of longer than typical duration between released reports.

---

[33] In this case, a false alarm does not indicate the accuracy of forecasting. This is because it indicates where conditions are predicted below the *alternate* minima, and observed above the *landing* minima, as such the forecast may be accurate although the *system* has produced the false alarm.

Check-sums for forecast duration

The total duration of each operational assessment was calculated within each forecast, creating an overall summary of the forecasting reliability. This information was compared to the calculated forecast validity period by summing the duration of each individual comparison segment within each non-overlapping forecast timeline (as discussed on page 78).

A table summarising the operational assessments within each forecast was produced, which also showed the resultant calculation of a check-sum. Table 31 shows an example of this process, using resolved operational assessment segments shown in Figure 17, with timing data calculated from Table 30.

Table 31: Sample summary of forecast reliability – summarised from Table 30

Forecast number	Assessment type	Total duration of assessed elements (minutes)				Validity duration	Check-sum (assessed elements minus validity)
		Hit	Miss	False alarm	Correct rejection		
1 (from Table 30)	Forecast alternate minima vs reported landing minima	75	15	30	600	720	0
2	…	…	…	…	…	…	…

Each forecast was condensed into a single row of information providing an overview of assessed conditions and the calculation of the check sum as follows:

$$Check\ sum = (Hit + Miss + False\ alarm + Correct\ rejection) - Validity\ duration$$

If the check sum was equal to zero, this provides an indication that no data is missing from the record, such as would occur if a forecast or report segment was not decoded correctly. A flag was placed against every record to confirm that the check-sum value was valid.

It is important to note that forecast validity times are determined at very early stages of the information decode, compared to the assessment criterion, which are the result of bringing together many calculations and assessments. Therefore, it was considered that the check sum would capture most timing anomalies inadvertently introduced and pre-existing in the data.

Duplicate and simultaneously released reports

In some cases, more than one METAR or SPECI report was released within the same minute of time. This meant that more than one record was used by the comparison algorithm for that duration to represent the conditions until the release of the next report (after that specific minute in time).

These records needed to be identified to prevent over-counting or over-representation of these times. This was particularly important for results and analysis based on likelihood, which used the aggregation of the time duration for assessment.

To alleviate duplicate time comparisons, a single METAR or SPECI was selected wherever more than one was present in the data within the same minute of time. The selection of the METAR or SPECI was based on following hierarchy:

- Most recent issue time – sometimes there was variation between the release or issue time and the published report time, then
- SPECI over METAR – SPECIs were always selected in preference over METARs, then
- The last report in the data provided (the physical positioning of the reports in the raw data)

Time exceeding typical duration between forecasts and reports

Analysis was performed to identify TAFs, and METARs or SPECIs with longer time gaps between each record. This was important for METARs and SPECIs because at the time of writing, actual conditions were approximated using reports from the time of release until the release of the next report. For TAFs this was important as the release time between TAFs was used to calculate the likelihood of retrieval of that particular TAF over time, as discussed in ‘Calculating the availability of forecast comparison states’ on page 86.

For METARs and SPECIs, time periods longer than 60 minutes, and TAFs with time periods longer than 7 hours were identified for further review and exclusion. Analyses were also performed showing the distribution of the typical length of time between METARs and SPECIs. At the time of writing, for Mildura, the vast majority of METARs and SPECIs were released within 35 minutes of the prior report.

Treatment of invalid records

In summary, validation checks performed on the data set were:

- check-sums comparing the original decoded forecast validity period with the resolved comparison segments
- the assessment of duplicate or simultaneously released reports
- identification of longer than typical duration between released reports
- holistic simulation time vector check-sums.

Records identified as invalid were flagged for removal for each forecast-observation report comparison. To ensure that valid records surrounding these were evaluated fairly, the relevant times that these segments would have been active were used in the analysis with no assessment, and were recorded as ‘Not decoded’.

-

Note that if start times and end times do not line up at the following stage of analysis, this was recorded as a discontinuity and excluded.

Additional benefits of identifying these records was finding un-anticipated errors and aspects of the data, allowing further refinement of these processes. These checks were determined based on testing some of the core assumptions of the data in this study, specifically – check-sums of all resolved final segments will add to the validity time of the forecast; there were no METARs or SPECIs released at the same time; the data set was complete with no gaps or delay between forecasts or reports.

Methodology of results and analysis - simulation of the operational effect of forecast retrievals

The following chapter provides an overview of a computer based simulation algorithm, which was specifically developed as an analysis tool of operational weather reliability. In addition, the chapter will cover the main concepts, techniques and metrics used to calculate and characterise the likelihood and potential impact of weather forecasting on aviation operations.

The simulation used pre-configured resolved comparisons of forecasts (TAFs and TTFs) with weather reports (METARs and SPECIs), as described in the section ‘Overall forecast assessments for operational reliability’ on page 77. Both forecasts and weather reports are linked to runway approach data selected according to specifications of nominated aircraft. The process of forming, combining and configuring this data to a nominated rule-set (legislative or otherwise) is specified in depth in the chapter Methodology of data extraction, formation and evaluation on page 7.

Note that all data joined together in the evaluation and extraction phases were linked to the following comparisons allowing analysis of all extracted elements.

Simulating the operational effect of forecast retrievals at specific times prior to arrival

One of the primary objectives of the study was to determine the likelihood of retrieving a forecast predicting conditions significantly above those observed, in particular, conditions observed and reported below published landing minima limits.

In addition, determining the potential impact that this may have on aircraft was also central to the objectives of this study. To address this, estimates of the potential aircraft unexpected holding time, and the probable number of aircraft affected were calculated. The determination of these concepts are discussed using a central case study to explain, with all figures (Figure 18, Figure 19 and Figure 20) and tables (Table 32 and Table 33) referring to the same fictitious set of data, made up of resolved TAF type forecasts. The process followed was identical for TTF type forecasts.

Analysis of individual resolved forecasts following forecast retrievals at nominated times prior to arrival

For weather related operational decision-making, it was considered important to understand the effects of varying forecast retrieval times prior to arrival. This is because operational effects may differ depending on the time of forecast retrieval, leading to different operational decisions, such as carrying additional fuel, or diverting.

Figure 18 shows how different operational results are possible for arrival at the same time (light blue diamonds labelled ‘AT 1’ and ‘AT 2’) when using separate, overlapping forecasts (TAF 1 and TAF 2). The two scenarios presented show forecast retrievals at different times (pink diamonds labelled ‘RT 1’ and ‘RT 2’), and different elapsed times from retrieval prior to arrival (labelled ‘e1’ and ‘e2’).

Forecast ‘assessment states’ are shown by the colour coding and outline of each resolved forecast comparison timeline, as described in the section ‘Formation of non-overlapping reliability timeline for each TAF and TTF’ on page 78, and illustrated in Figure 17.

Figure 18: Measuring operational effect of individual forecasts following retrieval at different times prior to arrival at the same time

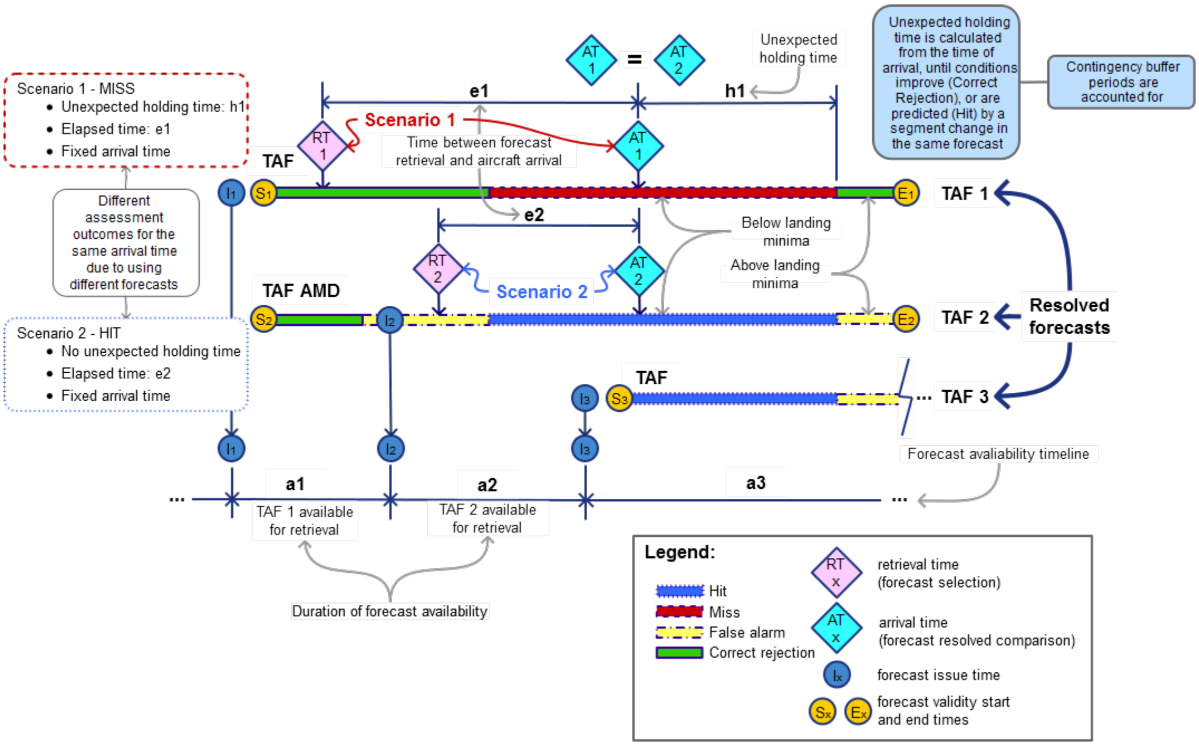

Forecast availability

In scenario 1, a ‘miss’ is recorded at time ‘AT 1’ because at the time of forecast retrieval, forecast ‘TAF 1’ was available. This is due to the retrieval time ‘RT 1’ falling between the ‘TAF 1’ issue time (blue circle labelled ‘I_1’), and the ‘TAF 2’ issue time (blue circle labelled ‘I_2’). This duration in time, labelled ‘a1’ refers to the availability of the forecast.

Conversely, for scenario 2, a ‘hit’ was recorded at the same time, because the forecast retrieval occurred when the amended TAF, ‘TAF 2’ was available. The duration of availability of TAF 2, labelled ‘a2’ spans between the TAF 2 and TAF 3 issue times (‘I_2’ and ‘I_3’).

The forecast availability is important when considering out of sequence forecast releases such as amended TAFs. Consider TAF 2, which covers the same validity period and amends TAF 1. The likelihood of a flight crew retrieving and using the information contained in TAF 1 is expected to be proportional to the availability period labelled ‘a1’.

If the period between the issue of two forecasts is short, the likelihood of forecast retrieval also diminishes, and conversely, if the period between forecast issues is longer, a relative increase in the likelihood of retrieval is expected.

This likelihood of exposure to forecasts is discussed in further depth in the section ‘Formation of non-overlapping timelines of resolved forecast comparisons based on fixed retrieval times prior to arrival’ on page 86.

Individual forecast characteristic analysis

Analysis of each forecast comparison timeline allows the assessment of characteristics of individual resolved forecasts comparisons, providing an overall context. For example, if a resolved forecast comparison contains a ‘Miss’, it may be of interest to know if any of the part of the forecast surrounding this period contains any predictions below the alternate minima, or none at all. This is covered in further depth in ‘Formation of non-overlapping reliability timeline for each TAF and TTF’ on page 78.

Unexpected holding time

A potential unexpected holding time labelled ‘h1’ applies to scenario 1. This is based on the expectation that a flight crew would be un-prepared if basing their decision making and planning around TAF 1, which predicts conditions above the alternate minima, however, on arrival, conditions were reported as unsuitable for landing (below the landing minima). The calculated holding time starts at the time of arrival and ceases effect after duration ‘h1’ when either conditions improve (which applies in this case), or a forecast segment in the same forecast would be expected to trigger contingency planning (conditions forecast below the alternate minima).

No unexpected holding time applies to TAF 2 as it is expected that the crew would be reasonably warned that a contingency plan was required.

Formation of non-overlapping timelines of resolved forecast comparisons based on fixed retrieval times prior to arrival

To gain a holistic understanding of operational reliability of forecasting, it was desired to know the historical availability of particular forecasting attributes, such as the expected number of ‘misses’ over a period of a day or month, or over all data available.

The likelihood of a forecast being retrieved is expected to be related to the duration of forecast availability described above, the number of aircraft arrivals at given times, and the times prior to arrival that flight crew or other operations personnel are likely to retrieve these forecasts.

The time a particular comparison scenario exists can be used to indicate the chances of exposure to particular conditions if an aircraft were to arrive at a particular time of day or month. Multiple timelines were produced in 15 minute retrieval time offsets, allowing the assessment of the relative likelihood of forecast reliability as the retrieval time prior to arrival increased.

Calculating the availability of forecast comparison states

While Figure 18 showed the characteristics of forecasting, it did not provide an indication of how likely a scenario like this would arise. Since forecasts are only available for discrete periods of time (the forecast availability discussed above), this represents the maximum envelope of time that the forecast could be retrieved.

For example, if a forecast was available for retrieval for 24 hours in one year, the odds of randomly selecting this forecast in one year would be approximately 1 in 365 days. Of course, the actual retrieval of forecasts is not expected to be random, which is accounted for in part by approximating the number of aircraft affected, discussed on page 89.

Figure 19 shows the process of creating a single timeline, simulating the effect of retrieving the available forecast (shown in Figure 18) for every minute of time at a given location, and recording the forecast-report comparison for that forecast at a constant time, labelled ‘e’ after forecast retrieval. For reference, Figure 19 refers to the same fictitious data as Figure 18.

Figure 19: Creating a non-overlapping forecast comparison timeline to assess operational effect of all forecast retrievals at the same time offset prior to arrival

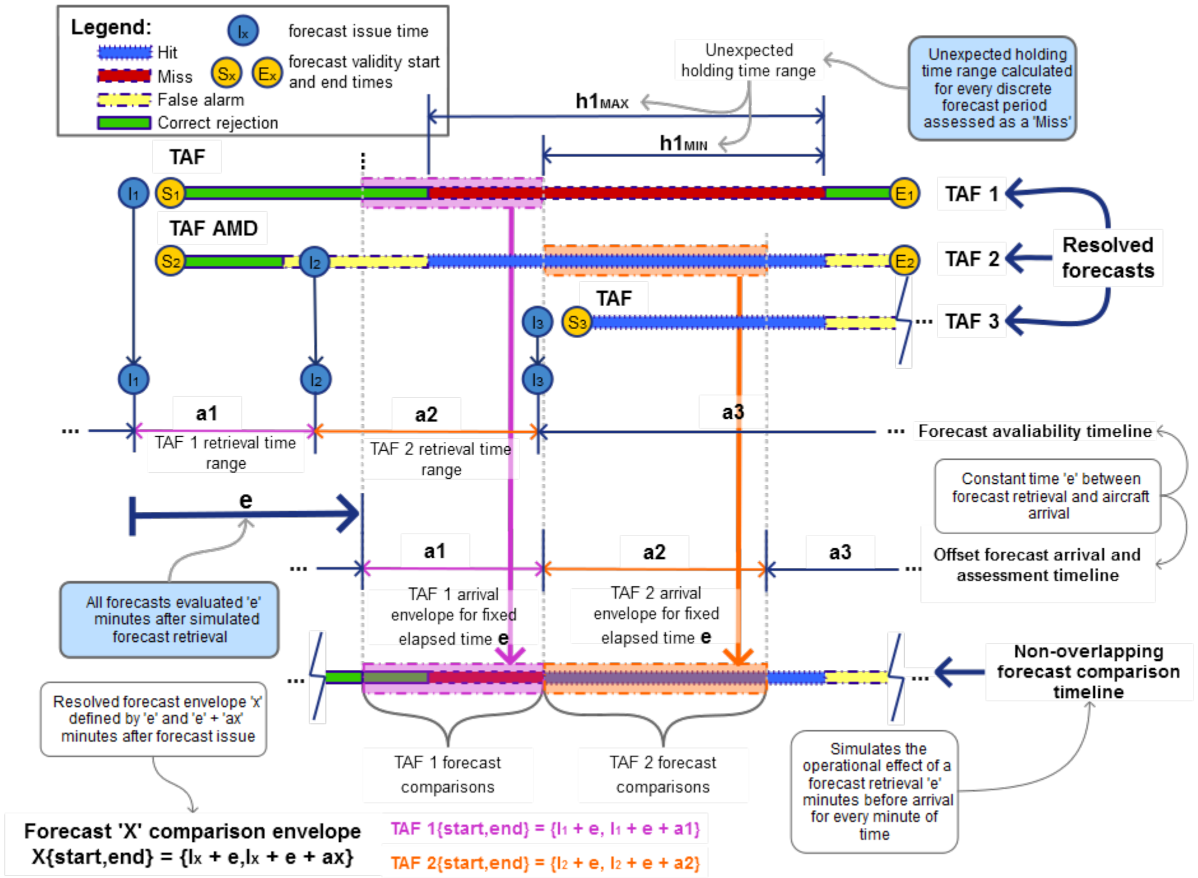

In this case, rather than having a single retrieval time point, as in Figure 18, a retrieval time *range* exists for each forecast, corresponding to the forecast availability. If the start and end times of the time range are shifted by the elapsed time ‘e’, an aircraft arrival time envelope can be calculated for every forecast.

The resolved forecast comparison data within this resultant time range is then inserted into the ‘non-overlapping forecast comparison timeline’, as labelled. For example, the assessment state on arrival TAF 1 started as a ‘Correct rejection’, followed by a period assessed as a ‘Miss’.

To illustrate this further, Table 32 contains sample calculations based on Figure 19, with a constant retrieval time prior to arrival set at 90 minutes.

Using TAF 1 as an example, it can be seen that the availability, or retrieval time range is recorded between the issue of TAF 1 at 0540, and the issue of TAF 2 at 0705, meaning that if a request for the current forecast was submitted during this time, the forecast TAF 1 would be retrieved. The time difference of 85 minutes corresponds to ‘a1’ in Figure 19.

The arrival envelope is calculated by shifting the availability period start and end times by 90 minutes from 0540 to 0710, and 0705 to 0835 respectively. This new arrival envelope corresponds to the only period in time that TAF 1 can have operational effect given a fixed 90 minute period between retrieval and arrival. In the case of TAF 1, a forecast retrieval in the first 30 minutes of availability would result in arrival during a ‘correct rejection’. Subsequently, for a constant 90 minute retrieval time offset, the availability of the forecast comparison state ‘Miss’ was 55 minutes for TAF 1.

Table 32: Sample calculations from Figure 19 of assessment states and holding time for forecast retrievals 90 minutes prior to arrival

Retrieval time prior to arrival (e) (minutes)	Forecast label	Issue time	TAF comparison start time	TAF comparison end time	Forecast retrieval availability period (a)	Holding duration		Availability by comparison states at arrival 90 minutes after forecast retrieval				
						h_{MIN}	h_{MAX}	Hit	Miss	False alarm	Correct rejection	Total
90	TAF 1	0540	0710	0835	85	135	190	0	55	0	30	85
	TAF 2	0705	0835	1030	115	0	0	115	0	0	0	115
	TAF 3	0900	1030	N/A	N/A	0	0	20	0	35	0	N/A

The formation of non-overlapping forecast comparison timelines enabled an overall analysis of the systemic performance of all forecasts of the same type at a location. For example, for the displayed period in Figure 19 and Table 32 the availability for arrival during conditions reported below the landing minima (recorded as a ‘hit’ or a ‘miss’), can be directly calculated from the timeline, in this case 55 (TAF 1) + 115 (TAF 2) + 20 (TAF 3) = 190 minutes.

Note that the availability of a forecast comparison state was not used to indicate the potential impact on aircraft safety. For example, the availability of the TAF 1 ‘Miss’ state is less than the overall assessed misses for TAF 1 because the amended forecast TAF 2 was released. For potential impact on aviation safety, an estimated holding time *range* was calculated.

Unexpected holding time range

The concept of unexpected holding time for a single time of arrival was introduced in Figure 18, as the time starting at the arrival during a forecast ‘miss’, until conditions below landing minima ceased or the forecast comparisons changed to a ‘hit’. Due to a timeline being produced (as opposed to a single arrival time), a *range* of holding times exist for every ‘miss’. This ranges from arrival at the start of the ‘miss’ up until the end, as shown in Figure 19. The minimum and maximum calculated holding times are used to provide an indication of the overall time an aircraft would be required to hold in the air.

Table 32 columns labelled ‘h_{MIN}’ and ‘h_{MAX}’ show values for calculated minimum and maximum holding times for TAF 1. This represents the shortest and longest durations of time that an aircraft would have to hold after retrieving TAF 1 at the start and end of the TAF 1 ‘Miss’ availability, as shown in the non-overlapping forecast comparison timeline of Figure 19. The minimum holding time corresponds to arriving at the end of the ‘Miss’ period, where conditions were predicted to degrade, or were reported to improve. Accordingly, the maximum holding time corresponds to arrival at the start of the ‘Miss’ period. Note the difference between the maximum and minimum holding times corresponds to the total availability of the ‘Miss’, in this case 55 minutes.

Note that multiple discrete periods of ‘misses’ either between forecasts or within forecasts (due to gaps between forecast ‘misses’) existed for some periods, which are calculated separately in each case. For the purposes of analysis, an additional ‘average’ holding time was calculated in each case. For example, this would equate to 162 minutes for TAF 1 of Figure 19 and Table 32.

Calculation of probable aircraft affected by merging flight activity data with forecast availability data

Calculating the availability of a forecast shows the times that a forecast would take effect given constant forecast retrieval prior to arrival times, as described from page 86. To indicate the potential effect this has had on actual flights, availability data was combined with an aircraft arrivals data model to estimate the number of aircraft exposed to each forecast comparison state.

Figure 20 illustrates this process showing the main three steps involved which are:

- Development of expected aircraft arrivals model
- Forming hourly splits in non-overlapping forecast comparison timelines
- Calculating expected aircraft arrivals for forecast comparison ‘states’.

Development of expected aircraft arrivals model

It was expected that during periods of weather where conditions were reported below the landing minima, aircraft arrivals would cease, or be significantly reduced during that period. To account for this a model of aircraft arrivals was produced using Airservices Australia flight activity data provided to the ATSB on an hour by hour basis, allowing relatively higher fidelity to calculate the total effect weather forecasting has on aircraft operations.

The objective of the model was to estimate the number of aircraft that *would* have arrived at a given time, had the weather been suitable for landing.

The algorithm developed allowed for different models of expected aircraft arrivals, however, at the time of writing, the arrivals data averaged in groups of the hour of day, weekend or weekday, and month. This data was then normalised by yearly arrivals, ensuring that the total expected arrivals in the model, matched the total yearly arrivals.

An example of an average group is the total of all aircraft arrivals at 10am, on weekends in June, a total of about 8 data points per year. The average of this group corresponds to the number of expected arrivals per hour for each group.

The aircraft flight activity data was grouped to a level that it was expected that flight arrivals would be similar. This data is expected to be conservative, as natural reductions in traffic movements due to poor weather will result in less aircraft being calculated as part of the total, then those actually affected.

A simple concept of this model is shown in Figure 20 with peak arrivals shown as three arrivals per hour between 9 and 10am, tapering by one aircraft per hour in the surrounding hours.

Forming hourly splits in non-overlapping forecast comparison timelines

Each non-overlapping forecast comparison timeline was adapted to summarise data by local time in hourly groups. This included the representation of unexpected holding time ranges and the availability to assessment states as shown in Figure 19.

Table 33 shows the results of this reclassification from that shown in Table 32. Note that this process no longer specifically references any forecasts, although the holding time is still calculated using specific forecast-report comparisons. For example, the estimated holding time still has a maximum recorded as 190 minutes, which reduces down to a minimum of 135 minutes, as shown in Table 32, however, additional divisions are made reflecting the 0800 split 20 minutes into the availability period of the ‘miss’. If more than one holding period existed within an hour, the average of each of the minimum and maximum values were used.

Availability periods were calculated by adding the time of each assessment state within each hour.

Calculating expected aircraft arrivals for forecast comparison ‘states’

The model of expected aircraft arrivals per hour was multiplied by the hourly proportion of each forecast comparison state in the same hour grouping according to the following equation:

$$\begin{aligned} &Aircraft\ affected \\ &\qquad = Hours\ of\ availability\ to\ assessment\ state\ \times Expected\ aircraft\ arrivals\ per\ hour \end{aligned}$$

The number of aircraft affected is expressed as a decimal rather than a whole number due to the averaging process used to determine the expected aircraft arrivals and the fractions of an hour that availability of comparison states occur (as the smallest unit of measurement in the analysis is one minute of time).

Once the number of aircraft affected by hour was determined, these figures were then aggregated as required. All aircraft affected calculations are performed at an hour by hour level throughout the study period. The results dataset also includes sunrise and sunset data to allow analysis relative to these times. However, it was decided to initially use local time to show alignment with forecast releases and flight schedules.

Table 33 shows the calculated aircraft affected for the data shown in Figure 20. For example, between 0900 and 1000, only a hit is recorded as available, therefore, this is calculated as 1.0 hour availability x 3 aircraft per hour = 3 aircraft. Another example is the calculation of a ‘Miss’ between 0700 and 0800 where 20/60 hours availability x 1 aircraft per hour = 0.33 aircraft affected.

Figure 20: Calculating the number of aircraft affected by different assessment states in each non-overlapping forecast comparison timeline (as shown in Figure 19)

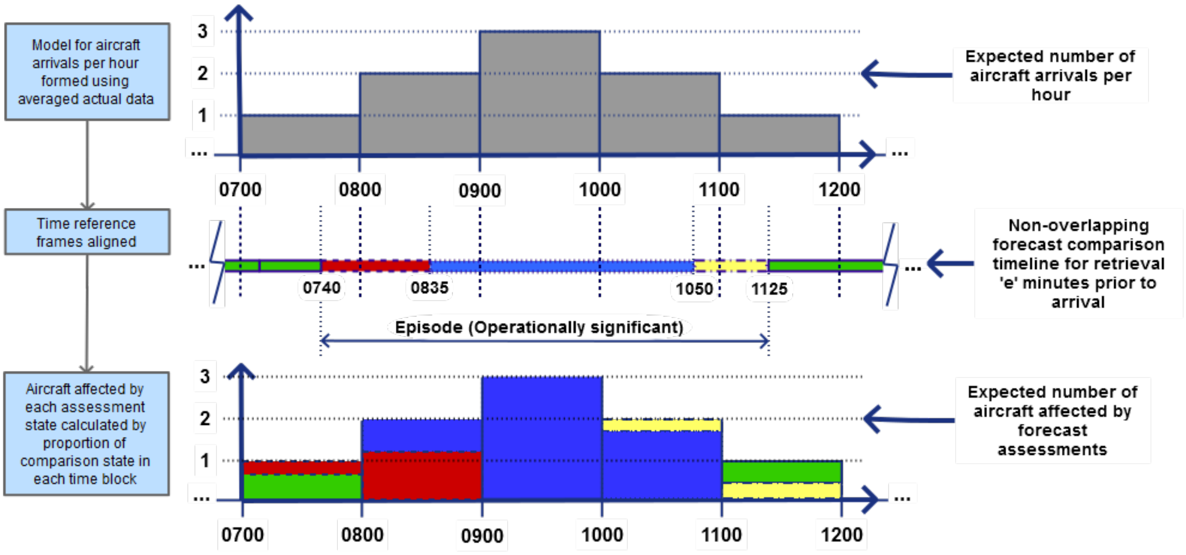

Table 33: Sample summary of operationally significant episode (shown in Figure 20) including number of aircraft affected, availability of assessment states, and estimated aircraft holding time by hour of day for forecast retrieval 90 minutes prior to arrival (shown in Table 32)

Hour of day	Estimated aircraft holding time (minutes)		Availability to assessment state (minutes)				Estimated affected aircraft			
	h_{MIN}	h_{MAX}	Hit	Miss	False alarm	Correct rejection	Hit	Miss	False alarm	Correct rejection
0700	170	190	0	20	0	40	0	0.33	0	0.67
0800	170	135	25	35	0	0	0.83	1.17	0	0
0900	0	0	60	0	0	0	3	0	0	0
1000	0	0	50	0	10	0	1.67	0	0.33	0
1100	0	0	0	0	25	35	0	0	0.42	0.58
Total	N/A	N/A	135	55	35	75	5.5	1.5	0.75	1.25

Grouped episodes

To evaluate periods with forecast or reported conditions below the alternate minima, an algorithm was developed to group non-overlapping forecast comparisons. Conditions were considered as falling within the same episode if the comparison state improved to a 'correct rejection' however returned to any other comparison state within one hour. A grouped episode is shown in Figure 20.

Combined assessment states used in simulation

Assessment states discussed thus far in the methodology have related to comparison of forecast alternate minima with reported landing minima, as shown in Figure 16 to indicate the possible operational effect. This shows the assessment 'state' of a forecast-report pair, providing an indication of both whether a selected aircraft model would be able to make a landing within limits, and whether a flight crew would be prepared for this scenario if using the forecast as the sole basis for operational decision making. This also discussed the calculation of other forecast and report comparisons 'states', specifically, comparing all combinations of forecast and reported alternate and landing minima.

A composite forecast-report comparison matrix was used in the forecast-report comparison analysis. This involved combining the forecast alternate vs reported landing minima and forecast alternate vs reported alternate minima comparison states, resulting in six logical comparison states[34], as shown in Figure 21. This corresponds to a cross-tabulation of the three possible reported conditions above the alternate minima, between the landing and alternate minima, and below the landing minima. Forecast conditions were unchanged above and below the alternate minima.

The output of the comparison was inserted into the results for use in the analysis. All metrics, such as number of aircraft affected and the likelihood of forecast comparison states were implemented in identical fashion to that described in 'Simulating the operational effect of forecast retrievals at specific times prior to arrival' on page 84. The calculation of aircraft holding time remained unchanged.

The new comparison matrix had no effect on the assessment of reported conditions below the landing minima as shown in Figure 16, but resulted in splitting the 'false alarm' and 'correct rejection' states into two groupings. The comparison matrix was considered to provide a 'fairer' measure of false alarms due to the separation of group '3a' in Figure 21, which contains reported conditions below the

---

[34] There were 16 theoretical combinations possible, eight of these were mutually exclusive and two were cases where the alternate minima was below the landing minima. These were used as a quality check when fusing these comparison matrices.

alternate minima. In addition, group '4a' also shows observations between the landing and alternate minima, where conditions were forecast above the alternate minima. These were previously classified as 'correct rejections', however indicate where the 'protections' of the difference between the landing and alternate minima acting as a systemic buffer has taken effect.

Note that the additional groupings only apply for the weather parameters visibility and ceiling, as the only values with differences between the landing and alternate minima thresholds.

Figure 21: Forecast predictions compared to actual observations including conditions between landing and alternate minima

		Reported observation		
		Below landing minima	**Between landing and alternate minima**	**Above alternate minima**
Forecast	**Below[35] alternate minima**	1 Landing not possible with contingency plan	3a Contingency plan required, landing possible	3b False alarm
	Above[35] alternate minima	2 Unsafe situation, landing not possible without contingency	4a Conditions not forecast, but landing possible	4b Correct rejection

The following summarises each comparison state:

1. Landing not possible with contingency plan: This is where a flight crew are using information indicating that an alternate minima is required, and the scenario eventuates where a landing is not possible at the planned destination.
2. Unsafe situation: This is where a flight crew receive information that indicates no contingency plan for an alternate aerodrome is required, however, it is not possible to land upon arrival at the destination.
3. Conditions forecast below the alternate minima, and observed above landing minima
    a. Weather conditions forecast and observed to be below alternate minima, with landing possible: This represents a scenario where a flight crew would have been able to land in degraded conditions below the alternate minima. Alternate minima criteria regarding flight planning would apply to these flights.
    b. False alarm: This situation is where this forecast would require alternate minima contingencies, however where conditions did not eventuate requiring this at the time of arrival.
4. Conditions forecast above the alternate minima, and observed above the landing minima
    a. Weather conditions forecast to be above the alternate minima, although observed to be below the alternate minima. Landing would be possible and alternate minima criteria would not apply. This situation represents a scenario where the inherent buffer between the alternate minima and landing minima comes into effect to account for forecast inaccuracy.
    b. Correct rejection: This situation is where conditions below alternate minima were not forecast or observed.

Assessment of forecasts predicting conditions below the landing minima

At the time of writing, forecast conditions between the alternate and landing minima, and below the landing minima had not been split into a comparison matrix for systemic analysis. However, analyses

---

[35] Below the minima means worse than the minima conditions. Above means better than the minima conditions.

were conducted comparing forecast visibility and ceiling both between the landing and alternate minima, and below the landing minima limits.

An assessment state was not developed initially because the focus of the initial study was on the operational reliability using the planning rule-set, and the assessment of the landing compared with alternate minima buffer. This is because forecast conditions below the landing minima do not impose the same operational planning requirement as conditions below the alternate minima, and therefore are not the primary focus for the evaluation of the planning rule-set.

The algorithm developed allows the implementation of these assessments, which would result in a 3 x 3 comparison state matrix, with the same reported observation divisions as shown in Figure 21, and new alternate minima divisions as discussed. Although this was not evaluated in the current analysis, it will be considered for future studies focused on the evaluation of forecast landing minima as a specific ‘decision threshold’.

Additional simulation functionality

Additional simulation functionality

The results of each simulation were stored in three separate databases relating to the summary, synchronisation and results of each simulation. The simulation was run from an SQL script invoking a series of stored procedures. A simulation specific configuration file was loaded in each case defining the parameters such as the date range, locations, forecast types, retrieval times prior to arrival and differences between the estimated time of arrival between secondary, and the primary location.

The simulation has been developed to run analyses with synchronised data, allowing analysis of multiple forecast retrievals in parallel. This allows the assessment decision making plans using multiple forecasts, and the calculation of the probability of concurrent forecasts assessed as ‘Misses’ in the dataset.

Due to the size of the data set and computing limitations, the data set could be progressively built, updated and inserted. This allowed ‘preliminary results’ to be established while waiting for final results to compile. This also assisted for development purposes, allowing testing and validation over all procedures using a small amount of the data. Any holes in existing data could be filled or additional data could be added at the beginning or end without having to re-produce the entire data set.

Assessment of scenarios across different times and locations

Coupled analysis over multiple forecasts and retrieval times

The non-overlapping forecast comparison timelines described on page 86 can be synchronised allowing the retrieval of multiple forecasts to be simulated at different retrieval time offsets. This used an algorithm developed to combine forecast comparison timelines into a single probable event sequence, allowing the simulation of decision making strategies based on the multiple sources of available information at different retrieval times. For example, this allows the simulation of a crew using a TAF at the time of planning and then retrieving a TTF near the point of last safe diversion to form an operational plan and take a particular course of action.

At the time of writing, all timelines were analysed separately to examine specific operational reliability of the forecasting at a constant time offset, and to separate this from the effects of flight crew being ‘passed’ the most current forecasts. It is important to note that at some point, an aircraft may be committed to travelling to a particular location from a certain time, and this analysis intends to also provide a reference of the relative likelihood of particular comparison states arising at constant retrieval time offsets.

Spatiotemporal analysis

Probable event sequences were also devised to assess the likelihood of recurrence of actual events or other possible scenarios involving more than one location, such as a planned destination and an alternate aerodrome. To ensure a realistic comparison across locations, forecasts could be synchronised on the relative times of the probable event sequence, rather than the reported time. This accounts for the estimated time of arrival (to differ depending on the time of diversion) for a planned destination and alternate destination.

For example, if a probable event sequence involves an aircraft arriving at an alternate location 20 minutes after the estimated arrival time at the planned destination, times are synchronised by subtracting 20 minutes from the alternate destination forecast timings and then aligning with the planned destination time.

Appendices

Appendix A – Sample data

Table 34: Regular expressions used to extract weather groupings from METAR and SPECI reports

<table>
<tr><th>Weather grouping</th><th>Primary regular expression</th><th>Notes</th></tr>
<tr><td>Wind group</td><td>(((\d{3}|VRB|///)(\d{2}|//)(/|G|)(\d{2}|)KT)|)</td><td></td></tr>
<tr><td>Visibility group</td><td>(((\d{4}($CompassPoints))\s{1,3}|(\d{4}|CAVOK|////)\s{1,3}){1,2})</td><td></td></tr>
<tr><td>RVR group</td><td>((R[0-3]?\d(L|C|R)?/P?M?\d{4}V?M?P?(\d{4})?(U|D|N)?\s{1,3}){1,3})</td><td></td></tr>
<tr><td>Weather phenomena group</td><td>((((VC|\+|-){0,2}(($WxDescriptor|$WxPhenomena){1,28}){1,20})\s{1,3}){1,20})</td><td></td></tr>
<tr><td>Cloud group</td><td>(((((FEW|SCT|BKN|OVC)\d{3}(TCU|CB)?|SKC|NSC|NCD)|([1-8][A-Z]{2,3}\d{3})|(VV\d{3}))\s{1,3}){1,20})</td><td></td></tr>
<tr><td>Temperature / Dewpoint group</td><td>((((MS?|-)?\d\d\.\d)/((MS?|-)?\d\d\.\d))\s{1,3})</td><td></td></tr>
<tr><td>QNH group</td><td>(Q?(\d{4}\.\d)\s{1,3})</td><td></td></tr>
<tr><td>Recent weather group</td><td>((RE((VC|\+|-){0,2}(($WxDescriptor|$WxPhenomena){1,28}){1,20})\s{1,3}){1,20})</td><td></td></tr>
<tr><td>Windshear group</td><td>((WS RW?Y? ?([0-3]\d(L|C|R)?)\s{1,3}){1,4})</td><td></td></tr>
<tr><td>Rainfall group</td><td>(RF(((\d\d\.\d)/(\d\d\d\.\d)\s)|((\d\d\.\d)/(\d\d\d\.\d)/(\d\d\d\.\d)\s)))</td><td>Extracted from within remarks section</td></tr>
<tr><td>Remarks and TTF groups</td><td>((RMK.*?)(TTF.*?)?)</td><td>Split into two separate columns of data</td></tr>
</table>

Table 35: Variables used in weather grouping extraction regular expressions

<table>
<tr><th>Variable name</th><th>Variable value</th></tr>
<tr><td>$CompassPoints</td><td>N|NE|E|SE|S|SW|W|NW</td></tr>
<tr><td>$WxDescriptor</td><td>MI|BC|PR|DR|BL|SH|FZ|TS</td></tr>
<tr><td>$WxPhenomena</td><td>DZ|RA|GR|SN|SG|DU|SA|SS|DS|GS|PL|FG|BR|FU|HZ|PO|SQ|FC|VA|IC</td></tr>
</table>

Figure 22: Data processes between delimited weather report data and extracted METAR table

METAR, SPECI and TTF delimited data table

Data map number	Column name	Sample text
1	BoMBulkMetarDataId	1044
2	DataImportId	1
3	YearDateTime	200901220830
4	BoMNotes	Error at line 13191.67: Invalid cloud group syntax
5	ReportTypeCode	S
6	ReportType	SPECIAWS
7	ICAOCode	YMIA
8	IssueTime	220830Z
9	BulkMetarData	23016/25KT 1000NE 2000 BLDU VV014 33.3/06.7 1001.4 RMK RF00.0/000.0 CLD:OVC011 VIS:2400 BV:13.5 IT:45.4 VER:2.1
10	WindDirection	230
11	AvgWindSpeed	16
12	MaxWindSpeed	25
13	Visibility	2000
14	MinVisibility	1000NE
15	RVRGroups	R16L/0650V0650D R16R/0250V0375N
16	WeatherPhenomenaGroups	BLDU TS VCSH
17	CloudGroups	VV014; 2ST006 7SC038 1CB032
18	Temperature	33.3
19	Dewpoint	6.7
20	QNH	1001.4
21	RecentWeatherGroups	RETS
22	WindshearGroups	WS RWY05
23	RainfallLast10minutes	00.0; 00.0
24	RainfallLast60minutes	NULL; 000.0
25	RainfallSince9am	000.0; 002.0
26	Remarks	RMK CLD:OVC011 VIS:2400 BV:13.5 IT:45.4 VER:2.1
27	TTFGroup	TTF:NOSIG

METAR decode table

Data from	Calculation type	Column name	Output text	Additional table(s)?
1	U	BoMBulkMetarDataId	1044	
	N/A	MetarDecodeId	26208	
3	U	YearDateTime	200901220830	
4	U	BoMNotes	Error at line 13191.67: Invalid cloud group syntax	
5	U	ReportTypeCode	S	
6	U	ReportType	SPECIAWS	
7	U	ICAOCode	YMIA	
8	U	IssueTime	220830Z	
3,8	D	ReleaseDateTimeUTC	2009-01-22 08:30:00 +00:00	
3,8	D	NextMETARDateTimeUTC	2009-01-22 09:00:00 +00:00	
9	U	OriginalReport	23016/25KT 1000NE 2000 BLDU VV014 33.3/06.7 1001.4 RMK RF00.0/000.0 CLD:OVC011 VIS:2400 BV:13.5 IT:45.4 VER:2.1	
10	U	WindDirection	230	
11	U	WindSpeedAvg	16	
12	U	WindSpeedMax	25	
13, 17	A	Visibility	2000	
13	E	CAVOK	0	
13, 17	A	VisibilityEstimateId	1	Yes
14	E	MinimumVisibility	1000	
14	E	VisMinDirection	NE	
14	D	VisMinDirection_degrees	45	
15	U	RVRGroups	R16L/0650V0650D R16R/0250V0375N	Yes
16	U	WeatherPhenomenaGroups	BLDU TS VCSH	Yes
16	D	PhenomenaDecode	Blowing dust and thunderstorms with showers in the vicinity	
17	E	VerticalVisibility	1400; NULL	
13, 17	A	CeilingEffective	1400; 3800	
13, 17	A	CeilingOctas	NULL; 8	
17	U	CloudGroups	VV014; 2ST006 7SC038 1CB032	Yes
17	E	SKC	0	
17	E	NSC	0	
17	E	NCD	0	
13, 17	A	CeilingEstimateId	1	Yes
18	U	Temperature	33.3	
19	U	DewPoint	6.7	
20	U	QNH	1001.4	
21	U	RecentWeatherGroups	RETS	Yes
21	D	RecentWeatherDecode	Recent thunderstorms	
22	U	WindshearGroups	WS RWY05	Yes
23	U	RFlast10mins	0.0; 0.0	
24	U	RFlast60mins	NULL; 0.0	
25	U	RFsince9am	0.0; 2.0	
26	U	Remarks	RMK CLD:OVC011 VIS:2400 BV:13.5 IT:45.4 VER:2.1	
26	E	SkyObsc	0	

Key - Calculation type	
U	Un-modified direct copy
D	Derived from data
A	Algorithm based
E	Extracted directly from field
N/A	Not applicable

Figure 23: Parent table data processes to convert delimited weather report data into extracted TTF table

METAR, SPECI and TTF delimited data table		
Data map number	**Column name**	**Sample text**
1	BoMBulkMetarDataId	91486
2	DataImportId	3
3	YearDateTime	200901061949
4	BoMNotes	NULL
5	ReportTypeCode	S
6	ReportType	SPECIAWS
7	ICAOCode	YPAD
8	IssueTime	061950Z
9	BulkMetarData	18008/10KT 6000 -DZ ... RMK ... TTF:FM1950 20010KT 9999 DZ BKN015 FM2230 20012KT 9999 BKN020
10	WindDirection	180
11	AvgWindSpeed	08
12	MaxWindSpeed	10
13	Visibility	6000
14	MinVisibility	NULL
15	RVRGroups	NULL
16	WeatherPhenomenaGroups	-DZ
17	CloudGroups	4ST013 7ST021
18	Temperature	17.3
19	Dewpoint	15.1
20	QNH	1011.6
21	RecentWeatherGroups	NULL
22	WindshearGroups	NULL
23	RainfallLast10minutes	00.0
24	RainfallLast60minutes	NULL
25	RainfallSince9am	000.0
26	Remarks	RMK CLD:SCT013 OVC021 VIS:9999 BV:13.5 IT:26.8 VER:2.1
27	TTFGroup	TTF:FM1950 20010KT 9999 DZ BKN015 FM2230 20012KT 9999 BKN020

Key - Calculation type	
U	Un-modified direct copy
D	Derived from data
A	Algorithm based
E	Extracted directly from field
N/A	Not applicable

Figure 24: Child table data processes to convert delimited weather report data into extracted TTF table

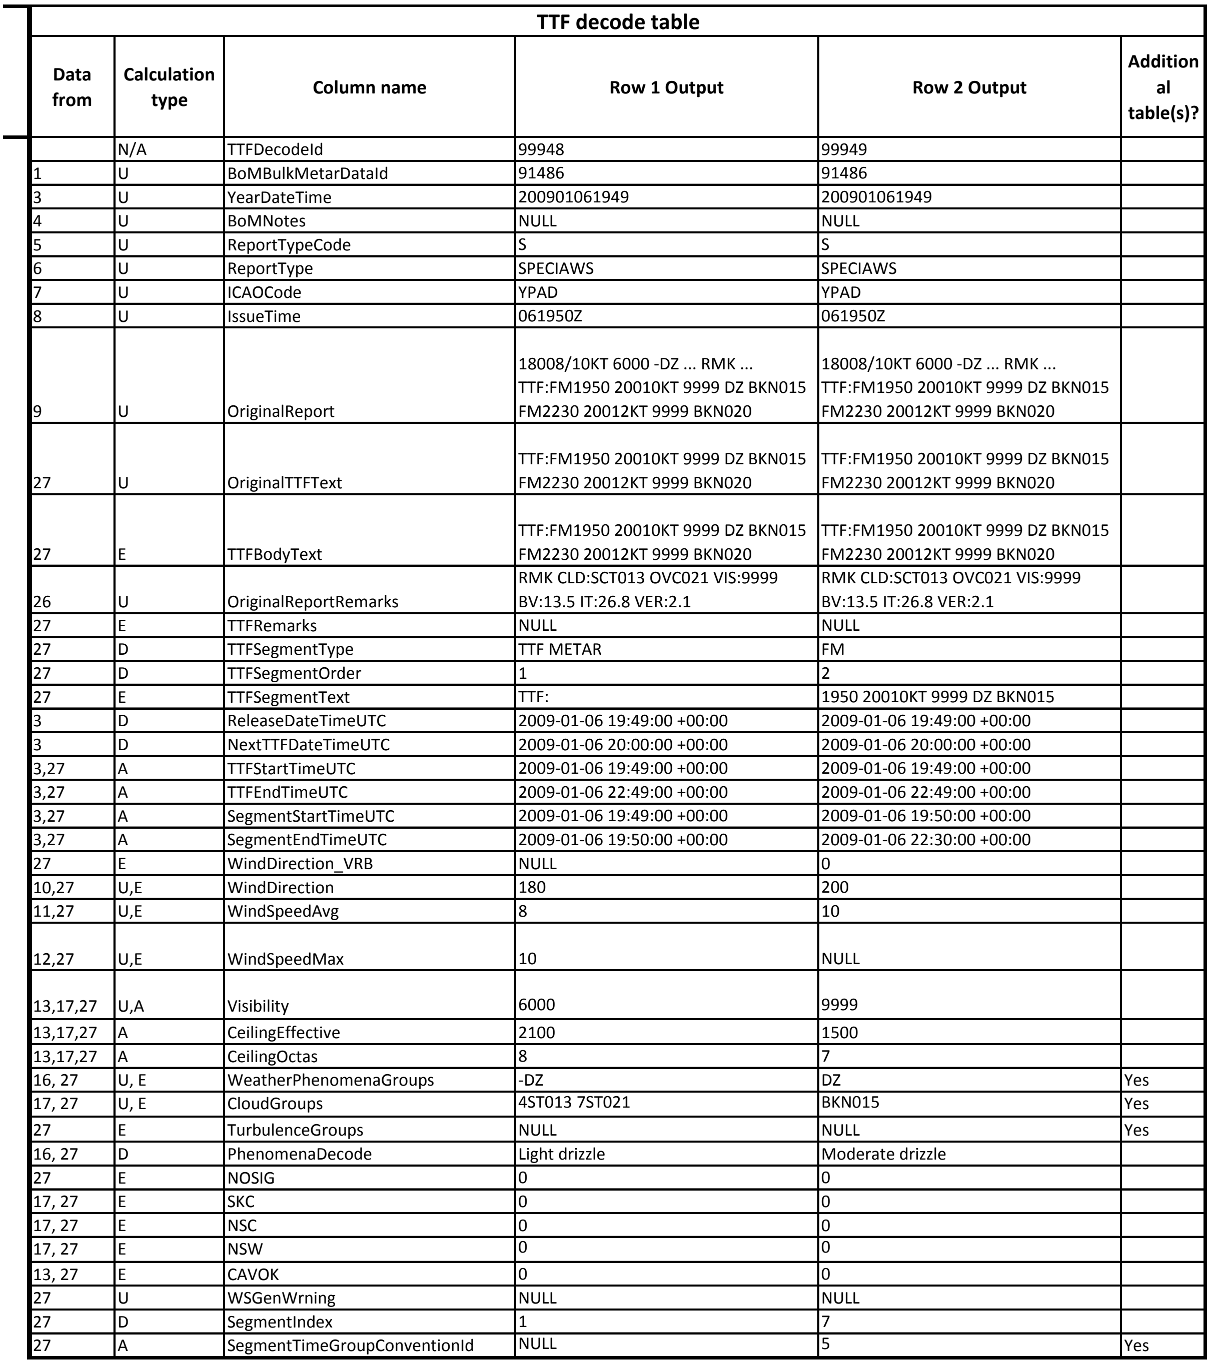

TTF decode table

Data from	Calculation type	Column name	Row 1 Output	Row 2 Output	Additional table(s)?
	N/A	TTFDecodeId	99948	99949	
1	U	BoMBulkMetarDataId	91486	91486	
3	U	YearDateTime	200901061949	200901061949	
4	U	BoMNotes	NULL	NULL	
5	U	ReportTypeCode	S	S	
6	U	ReportType	SPECIAWS	SPECIAWS	
7	U	ICAOCode	YPAD	YPAD	
8	U	IssueTime	061950Z	061950Z	
9	U	OriginalReport	18008/10KT 6000 -DZ ... RMK ... TTF:FM1950 20010KT 9999 DZ BKN015 FM2230 20012KT 9999 BKN020	18008/10KT 6000 -DZ ... RMK ... TTF:FM1950 20010KT 9999 DZ BKN015 FM2230 20012KT 9999 BKN020	
27	U	OriginalTTFText	TTF:FM1950 20010KT 9999 DZ BKN015 FM2230 20012KT 9999 BKN020	TTF:FM1950 20010KT 9999 DZ BKN015 FM2230 20012KT 9999 BKN020	
27	E	TTFBodyText	TTF:FM1950 20010KT 9999 DZ BKN015 FM2230 20012KT 9999 BKN020	TTF:FM1950 20010KT 9999 DZ BKN015 FM2230 20012KT 9999 BKN020	
26	U	OriginalReportRemarks	RMK CLD:SCT013 OVC021 VIS:9999 BV:13.5 IT:26.8 VER:2.1	RMK CLD:SCT013 OVC021 VIS:9999 BV:13.5 IT:26.8 VER:2.1	
27	E	TTFRemarks	NULL	NULL	
27	D	TTFSegmentType	TTF METAR	FM	
27	D	TTFSegmentOrder	1	2	
27	E	TTFSegmentText	TTF:	1950 20010KT 9999 DZ BKN015	
3	D	ReleaseDateTimeUTC	2009-01-06 19:49:00 +00:00	2009-01-06 19:49:00 +00:00	
3	D	NextTTFDateTimeUTC	2009-01-06 20:00:00 +00:00	2009-01-06 20:00:00 +00:00	
3,27	A	TTFStartTimeUTC	2009-01-06 19:49:00 +00:00	2009-01-06 19:49:00 +00:00	
3,27	A	TTFEndTimeUTC	2009-01-06 22:49:00 +00:00	2009-01-06 22:49:00 +00:00	
3,27	A	SegmentStartTimeUTC	2009-01-06 19:49:00 +00:00	2009-01-06 19:50:00 +00:00	
3,27	A	SegmentEndTimeUTC	2009-01-06 19:50:00 +00:00	2009-01-06 22:30:00 +00:00	
27	E	WindDirection_VRB	NULL	0	
10,27	U,E	WindDirection	180	200	
11,27	U,E	WindSpeedAvg	8	10	
12,27	U,E	WindSpeedMax	10	NULL	
13,17,27	U,A	Visibility	6000	9999	
13,17,27	A	CeilingEffective	2100	1500	
13,17,27	A	CeilingOctas	8	7	
16, 27	U, E	WeatherPhenomenaGroups	-DZ	DZ	Yes
17, 27	U, E	CloudGroups	4ST013 7ST021	BKN015	Yes
27	E	TurbulenceGroups	NULL	NULL	Yes
16, 27	D	PhenomenaDecode	Light drizzle	Moderate drizzle	
27	E	NOSIG	0	0	
17, 27	E	SKC	0	0	
17, 27	E	NSC	0	0	
17, 27	E	NSW	0	0	
13, 27	E	CAVOK	0	0	
27	U	WSGenWrning	NULL	NULL	
27	D	SegmentIndex	1	7	
27	A	SegmentTimeGroupConventionId	NULL	5	Yes

Figure 25: Parent table data processes to convert delimited TAF data into extracted TAF table

TAF delimited data table		
Data map number	**Column name**	**Sample text**
1	BoMBulkTafDataId	299
2	DataImportId	2
3	YearDateTime	20090310224619
4	BoMNotes	NULL
5	OtherBoMNotes	NULL
6	ReportType	TAF
7	ICAOCode	YMIA
8	IssueTime	102246Z
9	BulkTafData	1100 1112 05010KT 9999 -SHRA SCT040 SCT100 PROB30 TEMPO 1100/1112 VRB20G45KT 1000 TSRAGR BKN010 SCT040 RMK T 23 27 29 26 Q 1017 1015 1013 1014
10	TafDataICAO	1100/1112 05010KT 9999 -SHRA SCT040 SCT100 PROB30 TEMPO 1100/1112 VRB20G45KT 1000 TSRAGR BKN010 SCT040
11	TafRemarks	RMK T 23 27 29 26 Q 1017 1015 1013 1014

Figure 26: Child table data processes to convert delimited TAF data into extracted TAF table

Aircraft specification	Boeing 737-800	Median values - category D aircraft
AircraftSpecificationsId	14	38
Manufacturer	Boeing	Median
ModelCommonName	B737-800	D
ApproachCategoryId	4	4
MultiVariantDesignId	2	3
ILSCategoryId	4	4
MTOW_kg	79015	140160
MLW_kg	65317	111584
Range_nm	3115	4010
VS0_knots	Unknown	109
VS1g_knots	Unknown	Unknown
VAT_knots	0	0
VRef_knots	141	146
ActualNavigationPerformance_nm	Unknown	0.05
MLWClimbGradient_percent	Unknown	7.8
TODA_m	2500	2900
LDR_m	1380	1715
AvgCruiseSpeed_kmh	853	877
MaxDemXW_kts	33	33
MaxXW30mRwyDry_kts	27	27
MaxXW30mRwyWet_kts	16	16
TailwindMax_knots	10	10
NumberOfEngines	2	2
EngineManufacturer	CFMI	Not applicable
EngineModel	CFM 56-7B27	Not applicable

MaxThrust_1Engine_kN	121.4	208.28
Power_1Engine_kW	Unknown	Unknown
FuselageWidth_m	3.76	4.49
FuselageHeight_m	4.01	4.78
SeatingAisles	1	2
SeatsAbreast	6	7
MaxPax	162	237
Notes	Vref calculated assuming MLW ~65k with Full Flaps	Not applicable
References	Wikipedia, Qantas Operations Manual, Boeing.com,http://www.qantas.com.au/travel/airlines/aircraft-seat-map-boeing-738/global/en	Not applicable
SpecialAlternateCapable	1	1
GBAS	1	1
RNAV GNSS	1	1
RNAV RNP	1	1
ILS	1	1
LOC	1	1
LOC/DME	1	1
NDB	1	1
VOR	1	1
VOR/DME	1	1
NDB/DME	1	1

TAF decode table

Data from	Calculation type	Column name	Row 1 Output	Row 2 Output	Additional table(s)?
	N/A	TafDecodeId	479	480	
1	U	BoMBulkTafDataId	299	299	
3	U	YearDateTime	20090310224619	20090310224619	
4	U	BoMNotes	NULL	NULL	
5	U	OtherBoMNotes	NULL	NULL	
6	U	ForecastType	TAF	TAF	
7	U	ICAOCode	YMIA	YMIA	
8	U	IssueTime	102246Z	102246Z	
9	U	OriginalReport	1100 1112 05010KT 9999 -SHRA SCT040 SCT100 PROB30 TEMPO 1100/1112 VRB20G45KT 1000 TSRAGR BKN010 SCT040 RMK T 23 27 29 26 Q 1017 1015 1013 1014	1100 1112 05010KT 9999 -SHRA SCT040 SCT100 PROB30 TEMPO 1100/1112 VRB20G45KT 1000 TSRAGR BKN010 SCT040 RMK T 23 27 29 26 Q 1017 1015 1013 1014	
10	U	ICAOTafBodyText	1100/1112 05010KT 9999 -SHRA SCT040 SCT100 PROB30 TEMPO 1100/1112 VRB20G45KT 1000 TSRAGR BKN010 SCT040	1100/1112 05010KT 9999 -SHRA SCT040 SCT100 PROB30 TEMPO 1100/1112 VRB20G45KT 1000 TSRAGR BKN010 SCT040	
11	U	Remarks	RMK T 23 27 29 26 Q 1017 1015 1013 1014	RMK T 23 27 29 26 Q 1017 1015 1013 1014	
10	E	TAFSegmentType	TAF	PROB TEMPO	
10	D	TAFSegmentOrder	1	2	
10	E	TAFSegmentText	1100/1112 05010KT 9999 -SHRA SCT040 SCT100	1100/1112 VRB20G45KT 1000 TSRAGR BKN010 SCT040	
10	E	TAFSegmentProbability	NULL	30	
3	D	ReleaseDateTimeUTC	2009-03-10 22:46:19 +00:00	2009-03-10 22:46:19 +00:00	
3	D	NextTAFDateTimeUTC	2009-03-11 05:15:12 +00:00	2009-03-11 05:15:12 +00:00	
3,10	A	TAFStartTimeUTC	2009-03-11 00:00:00 +00:00	2009-03-11 00:00:00 +00:00	
3,10	A	TAFEndTimeUTC	2009-03-11 12:00:00 +00:00	2009-03-11 12:00:00 +00:00	
3,10	A	SegmentStartTimeUTC	2009-03-11 00:00:00 +00:00	2009-03-11 00:00:00 +00:00	
3,10	A	SegmentEndTimeUTC	2009-03-11 12:00:00 +00:00	2009-03-11 12:00:00 +00:00	
3,10	A	BECMGPeriodEndUTC	NULL	NULL	
10	E	WindDirection_VRB	0	1	
10	E	WindDirection	50	NULL	
10	E	WindSpeedAvg	10	20	
10	E	WindSpeedMax	NULL	45	
10	A	Visibility	9999	1000	
10	A	CeilingEffective	10000	1000	
10	A	CeilingOctas	8	7	
10	E	WeatherPhenomenaGroups	-SHRA	TSRAGR	Yes
10	E	CloudGroups	SCT040 SCT100	BKN010 SCT040	Yes
11	E	TemperatureGroups	T 23 27 29 26	T 23 27 29 26	Yes
11	E	QNHGroups	Q 1017 1015 1013 1014	Q 1017 1015 1013 1014	Yes
11	E	TurbulenceGroups			Yes
10	D	PhenomenaDecode	NULL	NULL	
10	E	SKC	0	0	
10	E	NSC	0	0	
10	E	NSW	0	0	
10	E	CAVOK	0	0	
11	E	WSGenWrning	NULL	NULL	
10	D	SegmentIndex	1	57	
10	A	TAFTimeGroupConventionId	1	1	Yes
10	A	SegmentTimeGroupConventionId	NULL	1	Yes

Table 36: Aircraft specifications used for analysis – Boeing 737-800 values shown

Table 37: Flat file for instrument approach procedures data structure formation

Mapped to table	Column name	Example 1	Example 2
Approach plate	RunwayName	1	1
	ICAOCode[36]	YBBN	YBBN
	ApproachPlateName	RNAV-P (RNP) RWY 01	RNAV-P (RNP) RWY 01
	ApproachPlateReference	BBNGN07-140	BBNGN07-140

---

[36] The ICAO (aerodrome location) code is also inserted in the aerodrome circling minima and aerodrome alternate minima tables for reference

	ApproachPlateDate	2014-08-21	2014-08-21
	InactiveApproach		
Approach	ApproachName	RNAV-P RNP RWY 01	RNAV-P RNP RWY 01
	ApproachNameLong	RNAV RNP RWY 01 Approach P	RNAV RNP RWY 01 Approach P
	ApproachTypeId	3	3
	CirclingAllowed	0	0
	Notes	MVD-4 Authorised CAT D Minima Only	MVD-4 Authorised CAT D Minima Only
Approach minima	MinimaCalculatedUsingForecastQNH		
	SpecialAltnMnmApplies	0	0
	ReqNavPerf	0.20	0.20
	MinMAPGrad_percent	5.40	5.40
	MinMAPGradToAlt	3000	3000
	ILSCategoryId	7	7
	MVDCategoryId	2	3
	MinDescentAlt_A		
	HeightAboveAerodromeRef_A		
	MinVisibility_m_A		
	RequiredRunwayVisualRange_A		
	MinDescentAlt_B		
	HeightAboveAerodromeRef_B		
	MinVisibility_m_B		
	RequiredRunwayVisualRange_B		
	MinDescentAlt_C	350	356
	HeightAboveAerodromeRef_C	338	344
	MinVisibility_m_C	900	900
	RequiredRunwayVisualRange_C		
	MinDescentAlt_D		356
	HeightAboveAerodromeRef_D		344
	MinVisibility_m_D		900
	RequiredRunwayVisualRange_D		
Aerodrome alternate minima	AlternateCeiling_FT_A		
	AlternateVisibility_m_A		
	AlternateCeiling_FT_B		
	AlternateVisibility_m_B		
	AlternateCeiling_FT_C	1237	1337
	AlternateVisibility_m_C	6000	7000
	AlternateCeiling_FT_D		1337
	AlternateVisibility_m_D		7000
	AlternateCeiling_FT_SPECIAL		
	AlternateVisibility_m_SPECIAL		
Aerodrome circling minima	CirclingMinDescentAlt_A		
	CirclingHeightAboveAerodromeRef_A		
	CirclingMinVisibility_m_A		
	CirclingMinDescentAlt_B		
	CirclingHeightAboveAerodromeRef_B		
	CirclingMinVisibility_m_B		
	CirclingMinDescentAlt_C		
	CirclingHeightAboveAerodromeRef_C		
	CirclingMinVisibility_m_C		
	CirclingMinDescentAlt_D		
	CirclingHeightAboveAerodromeRef_D		
	CirclingMinVisibility_m_D		

Figure 27: Approach plate example (partial data extraction shown in Table 37)

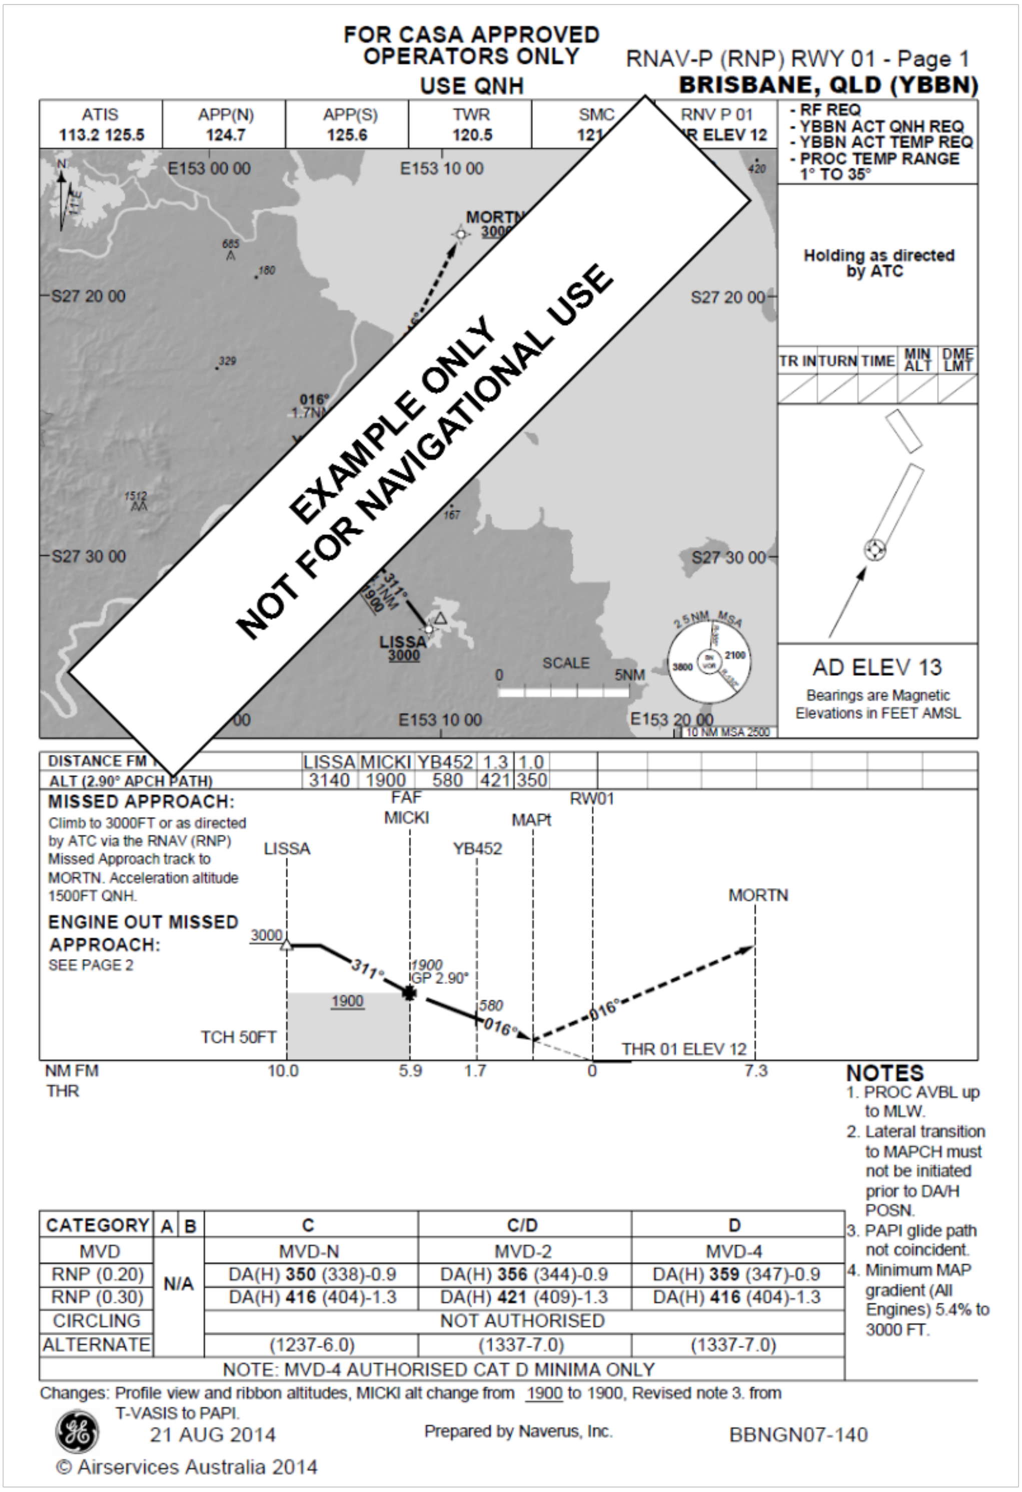

FOR CASA APPROVED
OPERATORS ONLY
USE QNH

RNAV-P (RNP) RWY 01 - Page 1
BRISBANE, QLD (YBBN)

ATIS	APP(N)	APP(S)	TWR	SMC	RNV P 01
113.2 125.5	124.7	125.6	120.5	121	R ELEV 12

- RF REQ
- YBBN ACT QNH REQ
- YBBN ACT TEMP REQ
- PROC TEMP RANGE 1° TO 35°

Holding as directed by ATC

TR IN	TURN	TIME	MIN ALT	DME LMT

AD ELEV 13

Bearings are Magnetic
Elevations in FEET AMSL

DISTANCE FM	LISSA	MICKI	YB452	1.3	1.0
ALT (2.90° APCH PATH)	3140	1900	580	421	350

MISSED APPROACH:
Climb to 3000FT or as directed by ATC via the RNAV (RNP) Missed Approach track to MORTN. Acceleration altitude 1500FT QNH.

ENGINE OUT MISSED APPROACH:
SEE PAGE 2

NOTES
1. PROC AVBL up to MLW.
2. Lateral transition to MAPCH must not be initiated prior to DA/H POSN.
3. PAPI glide path not coincident.
4. Minimum MAP gradient (All Engines) 5.4% to 3000 FT.

CATEGORY	A	B	C	C/D	D
MVD	N/A		MVD-N	MVD-2	MVD-4
RNP (0.20)			DA(H) **350** (338)-0.9	DA(H) **356** (344)-0.9	DA(H) **359** (347)-0.9
RNP (0.30)			DA(H) **416** (404)-1.3	DA(H) **421** (409)-1.3	DA(H) **416** (404)-1.3
CIRCLING			NOT AUTHORISED		
ALTERNATE			(1237-6.0)	(1337-7.0)	(1337-7.0)
NOTE: MVD-4 AUTHORISED CAT D MINIMA ONLY					

Changes: Profile view and ribbon altitudes, MICKI alt change from 1900 to 1900, Revised note 3. from T-VASIS to PAPI.

21 AUG 2014

Prepared by Naverus, Inc.

BBNGN07-140

Appendix B Lookup tables created for weather analysis

Table 38: Cloud type codes and names referenced in forecast and weather report tables

Cloud type Id	Cloud type name	Cloud type name long
1	CU	Cumulus
2	SC	Strato Cumulus
3	ST	Stratus
4	CI	Cirrus
5	AC	Alto Cirrus
6	CS	Cirro Stratus
7	AS	Alto Stratus
8	CB	Cumulonimbus
9	TCU	Towering Cumulus
10	CC	Cirro Cumulus
11	NS	Nimbo stratus

Table 39: Cloud coverage codes referenced in forecast and weather report tables

Cloud coverage Id	Coverage code	Coverage name
1	SKC	Sky clear
2	NSC	Nil significant cloud
3	NCD	Nil cloud detected
4	FEW	Few
5	SCT	Scattered
6	BKN	Broken
7	OVC	Overcast

Table 40: Weather intensity codes referenced in forecast and weather report tables

Weather intensity Id	Weather intensity name	Weather intensity symbol
1	Light	-
2	Moderate	
3	Heavy	+
4	Well developed	+
5	Not applicable	

Table 41: Weather vicinity codes references in forecast and weather report tables

Weather vicinity Id	Weather vicinity descriptor
0	Within 8km of aerodrome
1	8 to 16km from aerodrome

Table 42: Abbreviated time group convention list referenced in TAF and TTF type forecasts

Time Group Convention Id	Time Group Convention Name[37]
1	DDHH/DDHH
2	HHHH
3	DDHHHH
4	DDHHMM
5	HHMM
6	HHMM/HHMM

Table 43: Visibility estimate types used in METAR and SPECI reports

Visibility Estimate Id	Visibility Estimate
1	Reported in body text
2	Derived from vertical visibility
3	Remarks section visibility
4	Average of observations immediately prior and after METAR within set time limits
5	No estimation found - exclude record from analysis (Set visibility to 9999)

Table 44: Ceiling estimate types used in METAR and SPECI reports

Ceiling Estimate Id	Ceiling Estimate
1	Reported in body text
2	Cloud CLR BLW XXX derivation
3	Remarks section cloud groups
4	Slant Visibility Height where CLD: SKY MAY BE OBSC reported
5	Average of observations immediately prior and after METAR within set time limits
6	No estimation found - exclude record from analysis (Set ceiling to 99999)

Table 45: Aircraft performance categories

Approach Category Id	Approach Category	Minimum Approach Speed (knots)	Maximum Approach Speed (knots)
1	A	0	90
2	B	91	120
3	C	121	140
4	D	141	165
5	E	166	300
6	SPECIAL	NULL	NULL

Table 46: Instrument landing system (ILS) categories

ILS Category Id	ILS Category	Descision Height (feet)	Runway Visual Range (metres)	Visibility Minimum (metres)
1	I	200	550	800
2	II	100	300	NULL
3	IIIa	100	175	NULL

---

[37] Time group convention abbreviations: DD – day of month, HH – hour of day, MM – minute of hour

4	IIIb	50	50	NULL
5	IIIc	0	0	NULL
6	Not specified	NULL	NULL	NULL
7	Not applicable	NULL	NULL	NULL

Table 47: Landing approach types

Approach Type Id	**Approach name**
1	GBAS GNSS
2	RNAV GNSS
3	RNAV GNSS (RNP)
4	ILS
5	LOC
6	LOC/DME
7	NDB
8	VOR
9	VOR/DME
10	NDB/DME

Appendix C – Assessment of a safe forecast

AIP ENR 1.1 subsection 58 Working Summary used for algorithm development

Alternate aerodrome requirements

Suitable as a destination for the flight

Does not require an alternate

Requirements for an alternate to be set

- When forecast not available or provisional (58.1.3)
- Weather (58.2.1)
    - Cloud more than scattered below the alternate minimum, OR
    - Visibility less than the alternate minimum, OR
    - Visibility greater than the alternate minimum, with PROB Fog, Mist, Dust or any other visibility restricting phenomenon below the alternate minimum, OR
    - Wind – crosswind or downwind component more than the maximum for the aircraft
- TAF Timing alterations
    - When weather at the destination are predicted to be below alternate minimum but forecast to improve above at a specific time – can proceed if 30 minutes fuel is held beyond that time (58.2.3)
    - When FM or BECMG segments predict operational conditions these become effective 30 minutes before the published start time (58.2.7a)
    - When FM or BECMG segments predict the removal of operational conditions, this segments becomes effective 30 minutes after the published FM start time or 30 minutes after the published BECMG end time (58.2.7b).
    - 30 minutes is to be added to the beginning and end of started INTER and TEMPO segments predicting thunderstorms, associated severe turbulence or any conditions placing a requirement for an alternate aerodrome (58.2.8)
- TTF Timing alterations
    - When weather at the destination are predicted to be below alternate minimum but forecast to improve above at a specific time – can proceed if 30 minutes fuel is held beyond that time (58.2.3)
    - No other timing buffers applied (58.2.9)
- TTF and TAF Interpretation according to AIP
    - Sole reference to TTF allowed (58.2.9)
    - Max visibility used for TTF assessment of alternate, even when two visibilities included (58.2.10)
    - Flights which cannot use TTF will use TAF until such time as the TTF validity covers the ETA. (58.2.11)
    - Alternate minima used are those published on instrument approach charts (58.2.12a). Provisional or unavailable forecasts require an alternate (safe outcome) (58.2.12b). All aerodromes in this study have at least 1 published instrument approach procedure, and therefore an alternate of lowest safe altitude +500 feet and visibility greater than 8 kilometres is not applied (58.2.12c).

Appendix D – Development of statistical tests in SQL

Mann-Whitney-Wilcoxon ranked test

This test was developed to compare the ranking distributions between sets of ordinal or continuous weather parameters. A non-parametric test was required due to most data having a non-normal distribution. This can be used in conjunction with others tests to give an indication of significant differences between median values.

Method:

Numeric ranks were assigned to all observations within a set of data, for example observed visibility. This data was grouped into a binary category, for example a set of observed visibilities could be grouped by whether or not the particular moment in time could be evaluated or not to give an indication of the level of possible variation in the evaluated dataset.