

TIDB2

Implementation and Testing Object Extensions
of Open Source RDBS for Meteorological Data

João Simões - IM, Portugal

Maria Monteiro – IM, Portugal

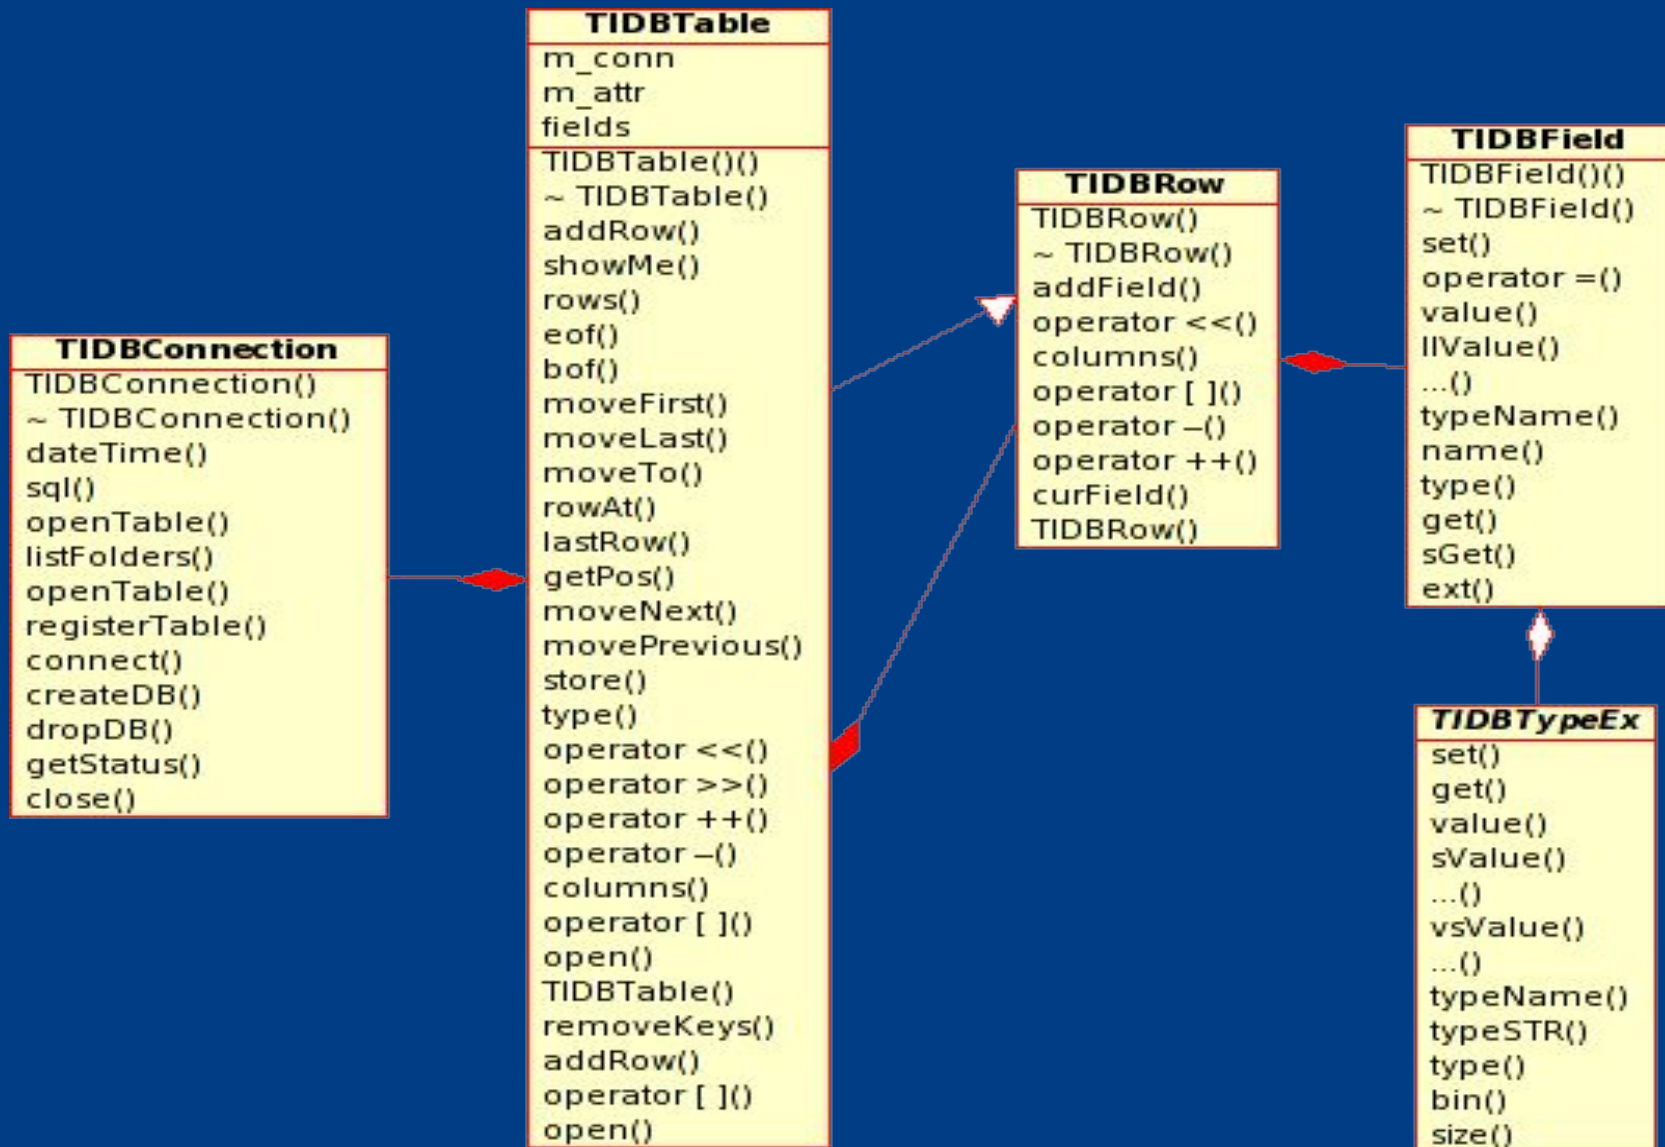
António Amorim – FCUL, Portugal

10th ECMWF Workshop on
Meteorological Operational Systems
November 2005

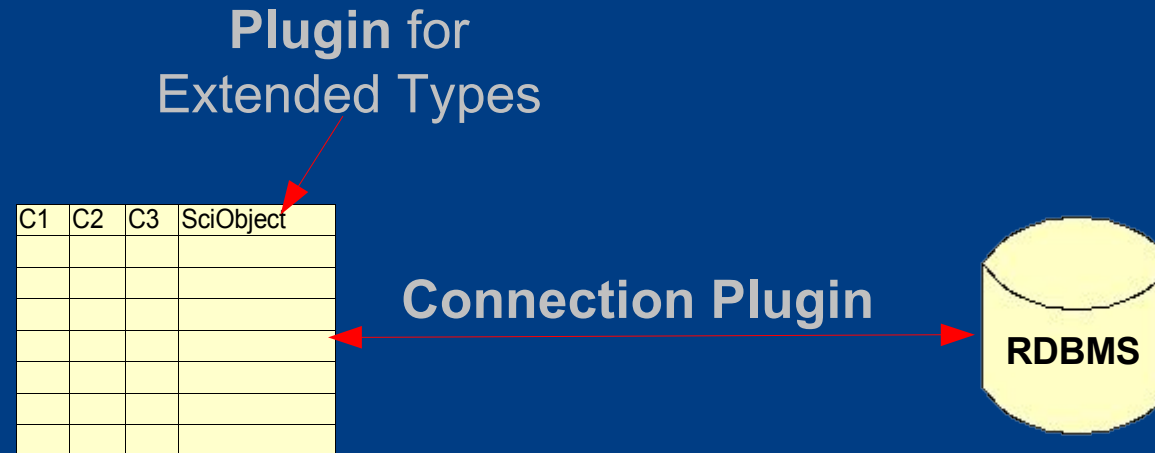
What is TIDB2?

- Fully featured temporal database.
- Simple and intuitive C++ interface.
- RDBMS independent (via runtime plugin).
- Oriented to store any kind of meteorological object or any scientific object in general (via runtime plugin).
- Automatic “index” creation, based on object's schema.
- Provides a set of simple tools for storing and retrieving objects.

General Structure

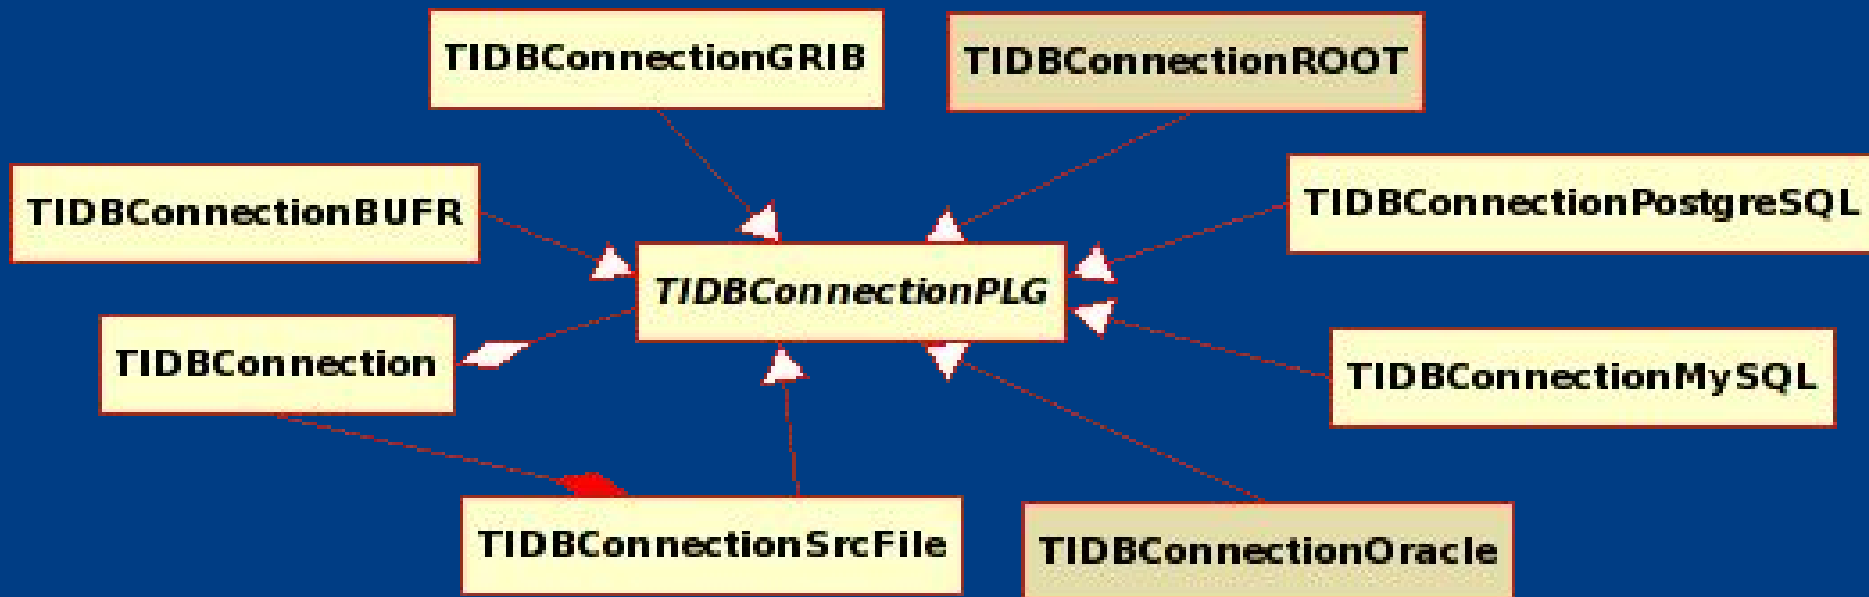


Plugin Architecture



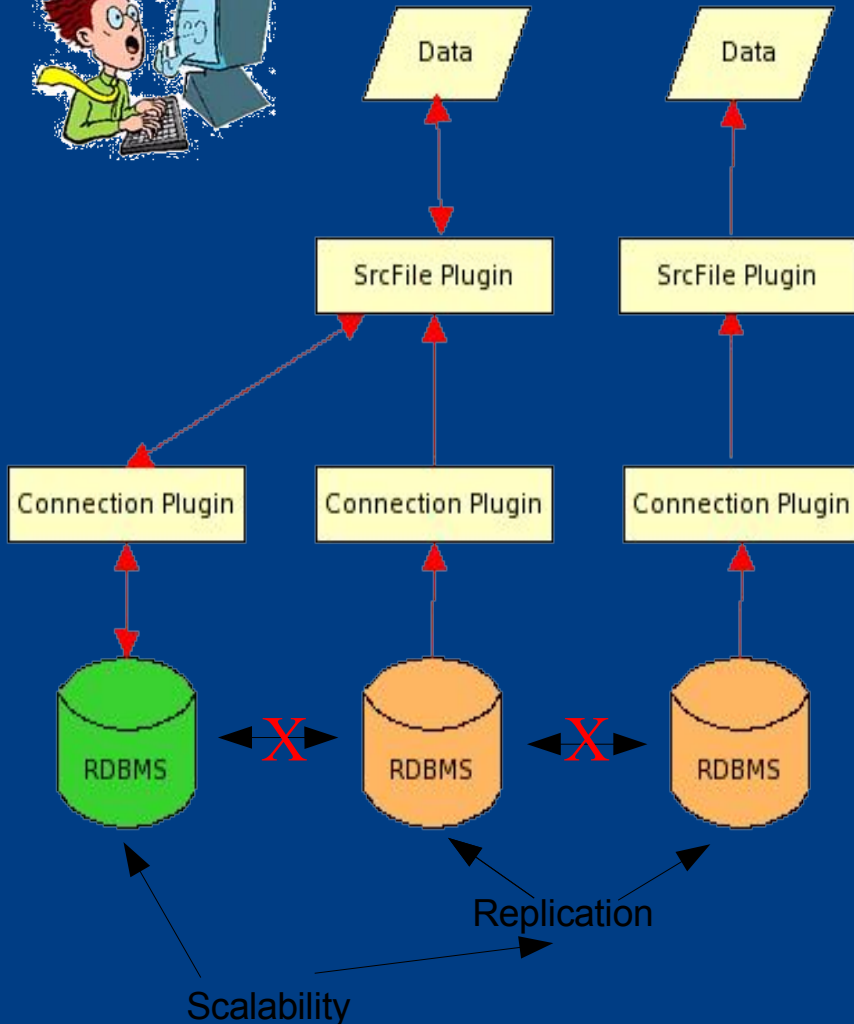
- Plugins are shared libraries loaded at runtime.
- Connection plugin stores and retrieves data from the DB.
- Extended Type plugin manages the columns containing scientific objects.

The Connection Object



- The TIDBConnection selects the appropriate plugin that will handle the connection (ex. mysql://, oracle://).
- All plugins implement TIDBConnectionPLG (providing all core functions to manipulate the database).

The Source File Plugin



- Uses a Debian “apt-get” like mechanism.
- Servers references are written to a source list file:
 - DB/Connection/Time Period.
- Makes scalability simple.
- Makes replication simple.

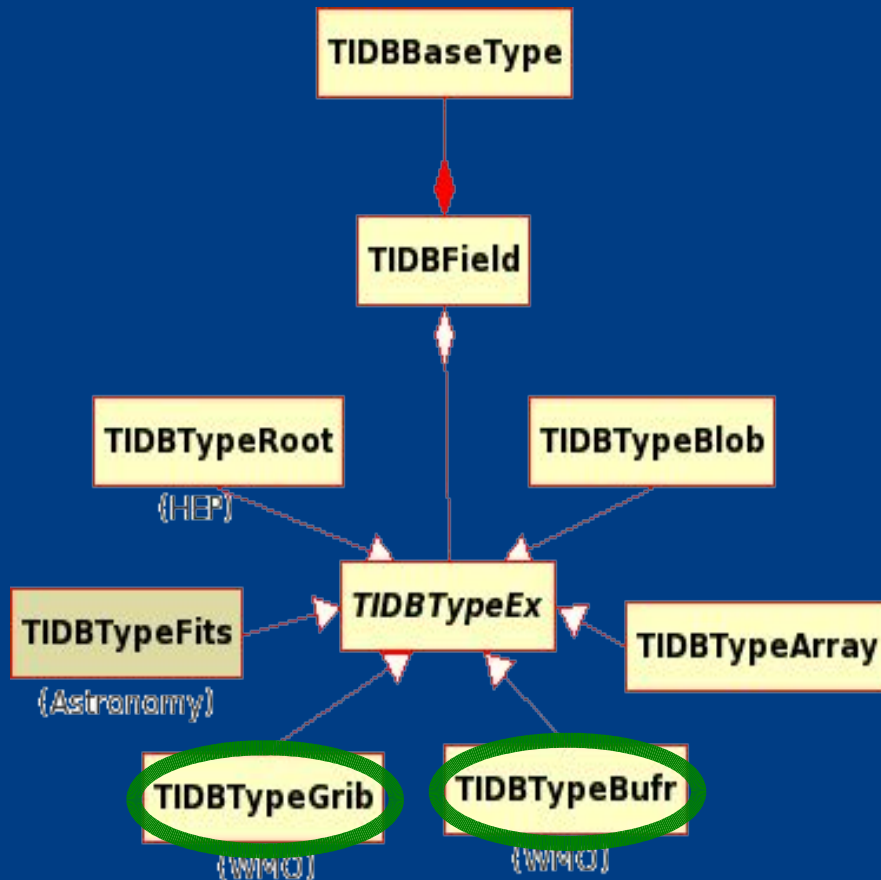
The TIDB2 Transient Table

TIDBTable is available when returning result sets and for storing procedures:



- Tables can be built from scratch using a row as a model, filled with rows and stored.
- Tables can be retrieved from a DB by a TIDBConnection, appended with rows and stored.
- The resulting table could be the result of a SQL query.
- Any external table can be registered in a the TIDB database, and opened as a TIDBTable.

The TIDBField and TIDBTypeEX



- TIDBField manages the data types.
- TIDBField implements the basic types.
- TIDBField provides an interface between the user and the extended types.
- With the appropriate plugin any data type can be supported.
- It's easy to fill a TIDBField with data.

The Special << Operator

```
TIDBRow MyRow(table) << 1 << "2" << 3.0;
```

- The “clever <<” operator automatically casts the data to the respective column type.
- This operator has a special behavior while streaming extended data types.
- TIDBRows can be streamed sequentially into a table.

Three Alternatives:

Atomized «complex data type» storage:

- The BLOB is splited into all it's elements.
- Lots of data redundancy or associations.
- Occupies a lot of storage space.

The data is kepted as BLOBs:

- Unsuitable for seeking objects.
- Makes it impossible to quickly find the most relevant data properties.

Mixed mode TIDB2 approach.

Atomized

A	B	C	D	E	F	G	H	I	J	K	L
A	B	C	D	E	F	G	H	I	J	K	L
A	B	C	D	E	F	G	H	I	J	K	L
A	B	C	D	E	F	G	H	I	J	K	L
A	B	C	D	E	F	G	H	I	J	K	L
A	B	C	D	E	F	G	H	I	J	K	L
A	B	C	D	E	F	G	H	I	J	K	L
A	B	C	D	E	F	G	H	I	J	K	L


BLOB storage

BLOB
BLOB
BLOB
BLOB
BLOB

TIDB2 mixed mode

A	E	H	BLOB
A	E	H	BLOB
A	E	H	BLOB
A	E	H	BLOB
A	E	H	BLOB
A	E	H	BLOB
A	E	H	BLOB
A	E	H	BLOB

```
TIDBRow MyRow(table) << SciObject;
```



- “Streamer” analyzes the internal structure of complex objects.
- “Streamer” automatically fills fields matching data inside the SciObject.
- The key fields are tunable (depends on the user defined table structure).

GRIB and BUFR Extended Types

- Any GRIB/BUFR data inside a GRIB/BUFR file could be loaded into a TIDBField - just need to be assigned to a filename.
- The TIDBField class provides methods to accede all GRIB/BUFR data and headers.
- BUFR tables should be stored in “/usr/local/etc/emos/bufrtables” or it's location should be set on a environment variable.
- Any BUFR/GRIB could be stored in a TIDBRow, making use of the “special streamer” that automates the index creation.

GRIB and BUFR tools

- Some tools have been developed, they use an environment variable with a connection string to connect to a TIDB2 database.
- storegrib, storebufr – for storing data inside a database.
- showgrib, showbufr – for retrieving a set of GRIBs/BUFRs from the database (search by center, a time interval and “SQL where clause”).
- Very useful when combined with scripting.
- Already integrated with Metview.

- Allow to store a link to GRIB or BUFR files instead of the objects themselves.
- Build a API for use with Fortran, making available various access methods:
 - C++
 - Fortran
 - Shell tools
 - KTIDBExplorer browser

The TIDB2 Browser: KTIDBExplorer

KonDBExplorer

File Edit Window Help

Name /ECMF/2005

	Since	Till	Para	Tir	T	Tab	LevelType	Level	La	Lor	LatMi	LatM	LonM	LonM	GribObj
12	2005-10-15/07:00:00	2005-10-15/07:00:00	130	0	0	128	100	1000.0	-1	1	75	15	-85	30	<GRIB DAT
13	2005-10-15/13:00:00	2005-10-15/13:00:00	130	0	0	128	100	1000.0	-1	1	75	15	-85	30	<GRIB DAT
14	2005-10-15/19:00:00	2005-10-15/19:00:00	130	0	0	128	100	1000.0	-1	1	75	15	-85	30	<GRIB DAT
15	2005-10-16/01:00:00	2005-10-16/01:00:00	130	0	0	128	100	1000.0	-1	1	75	15	-85	30	<GRIB DAT
16			130	0	0										<GRIB DAT
17			130	0	0										<GRIB DAT
18			130	0	0										<GRIB DAT
19			130	0	0										<GRIB DAT
20			130	0	0										<GRIB DAT
21			130	0	0										<GRIB DAT
22			130	0	0										<GRIB DAT
23			130	0	0										<GRIB DAT
24			130	0	0										<GRIB DAT
25			130	0	0										<GRIB DAT
26			130	0	0										<GRIB DAT
27			130	0	0										<GRIB DAT
28			130	0	0										<GRIB DAT
29			130	0	0										<GRIB DAT
30			130	0	0										<GRIB DAT
31			130	0	0										<GRIB DAT
32			130	0	0										<GRIB DAT
33			130	0	0										<GRIB DAT
34			130	0	0										<GRIB DAT
35			130	0	0										<GRIB DAT
36			130	0	0										<GRIB DAT
37			130	0	0										<GRIB DAT
38			130	0	0										<GRIB DAT
39			130	0	0										<GRIB DAT
40			130	0	0										<GRIB DAT
41			130	0	0	128	100	1000.0	-1	1	75	15	-85	30	<GRIB DAT
42			130	0	0	128	100	1000.0	-1	1	75	15	-85	30	<GRIB DAT
43			130	0	0	128	100	1000.0	-1	1	75	15	-85	30	<GRIB DAT
44			130	0	0	128	100	1000.0	-1	1	75	15	-85	30	<GRIB DAT
45			130	0	0	128	100	1000.0	-1	1	75	15	-85	30	<GRIB DAT

Filter options

Time Filter

Show all data

Show data in interval:

Since: 2005-10-14 13:51:32

Till: 2005-11-11 13:51:32

Field Filter

	Name	Oper	Value	x
1	ParamID	=	130	x
2	Time1	=	0	x
3	Time2	=		
4	TableVer	=		
5	LevelType	=		
6	Level	=	1000	x

Ok Cancel

The TIDB2 Browser: KTIDBExplorer

The screenshot shows the KondBExplorer application window. The interface includes a menu bar (File, Edit, Window, Help), a toolbar, and a tree view on the left. The tree view shows a hierarchy of folders: bufr (local, /tmp/bufr2, bufr.new, /tmp/bufrs), grib (local), and mysql (180.180.7.7, bufr, 1, 2005, grib, ALAD, ECMF, 2005, localhost, vpndb, fdb, ECMF, 2005, fdb2). The main area displays a data table for the path /1/2005. The table has columns: Since, Ce, Ty, St, Wn, Wmi, Height, Dry, Latitud, Longitu, and BufrObj. The data rows show various meteorological observations for the date 2005-11-05. A 'Filter options' dialog is open, showing a 'Time Filter' section with 'Show data in interval' selected. The 'Since' date is 2005-11-04 00:00:00 and the 'Till' date is 2005-11-12 00:00:00. The 'Field Filter' section shows a table with columns Name, Oper, Value, and a checkbox. The 'WmoBlock' field is selected with an equals sign operator and the value '8'. A 'Bufr detail' dialog is also open, showing a table of properties and values for the selected BufrObj.

Since	Ce	Ty	St	Wn	Wmi	Height	Dry	Latitud	Longitu	BufrObj
1	2005-11-05/00:00:00	98	0	1	8	501	29.0	293.4	39.45	-31.13 <BUFR DATA>
2	2005-11-05/00:00:00	98	0	1	8	509	54.0	290.4	38.77	-27.1 <BUFR DATA>
3	2005-11-05/00:00:00	98	0	1	8	512	72.0	291.0	37.73	-25.7 <BUFR DATA>
4	2005-11-05/00:00:00	98	0	1	8	515	100.0	288.0	36.97	-25.17 <BUFR DATA>
5	2005-11-05/00:00:00	98	0	1	8	524	82.0	291.3	33.07	-9.4 <BUFR DATA>
6	2005-11-05/00:00:00	98	0	1	8	531	32.0	289.0	39.35	-9.33 <BUFR DATA>
7	2005-11-05/00:00:00	98	0	1	8	533	26.0	284.8	37.0	-8.95 <BUFR DATA>
8	2005-11-05/00:00:00	98	0	1	8	541	103.0	284.2	37.95	-9.03 <BUFR DATA>
9	2005-11-05/00:00:00	98	0	1	8	543	18.0	280.8	41.7	-8.8 <BUFR DATA>
10	2005-11-05/00:00:00	98	0	1	8	545	77.0	282.5	41.23	-9.15 <BUFR DATA>
11	2005-11-05/00:00:00	98	0	1	8	548	179.0	281.8	40.15	-8.88 <BUFR DATA>
12	2005-11-05/00:00:00	98	0	1	8	554	8.0	286.5	37.02	-8.87 <BUFR DATA>
13	2005-11-05/00:00:00	98	0	1	8	567	562.0	278.0	41.27	-8.8 <BUFR DATA>
14	2005-11-05/00:00:00	98	0	1	8	568	1388.0	274.0	40.42	-8.8 <BUFR DATA>
15	2005-11-05/00:00:00	98	0	1	8	570	384.0	283.0	39.83	-8.8 <BUFR DATA>
16	2005-11-05/00:00:00	98	0	1	8	571	590.0	281.8	39.28	-8.8 <BUFR DATA>
17	2005-11-05/00:00:00	98	0	1	8	575	692.0	276.3	41.8	-8.8 <BUFR DATA>
18	2005-11-05/00:00:00	98	0	1	8	579	105.0	286.3	38.77	-8.8 <BUFR DATA>
19	2005-11-05/00:00:00	98	0	1	8	521	49.0	291.5	32.68	-8.8 <BUFR DATA>
20	2005-11-05/00:00:00	98	0	1	8	560	644.0	279.3	40.72	-8.8 <BUFR DATA>
21	2005-11-05/00:00:00	98	0	1	8	562	247.0	283.2	38.02	-8.8 <BUFR DATA>
22	2005-11-05/00:00:00	98	0	1	8	509	54.0	292.4	38.77	-8.8 <BUFR DATA>
23	2005-11-05/00:00:00	98	0	1	8	512	72.0	291.4	37.73	-8.8 <BUFR DATA>
24	2005-11-05/00:00:00	98	0	1	8	515	100.0	290.2	36.97	-8.8 <BUFR DATA>
25	2005-11-05/00:00:00	98	0	1	8	521	49.0	291.0	32.68	-8.8 <BUFR DATA>
26	2005-11-05/00:00:00	98	0	1	8	522	56.0	290.9	32.63	-8.8 <BUFR DATA>
27	2005-11-05/00:00:00	98	0	1	8	531	32.0	285.6	39.35	-8.8 <BUFR DATA>
28	2005-11-05/00:00:00	98	0	1	8	532	130.0	281.0	38.83	-8.8 <BUFR DATA>
29	2005-11-05/00:00:00	98	0	1	8	533	26.0	285.0	37.0	-8.8 <BUFR DATA>
30	2005-11-05/00:00:00	98	0	1	8	534	11.0	285.8	38.7	-8.8 <BUFR DATA>
31	2005-11-05/00:00:00	98	0	1	8	535	95.0	285.9	38.72	-8.8 <BUFR DATA>
32	2005-11-05/00:00:00	98	0	1	8	540	54.0	280.0	39.83	-8.8 <BUFR DATA>
33	2005-11-05/06:00:00	98	0	1	8	541	103.0	283.2	37.95	-8.8 <BUFR DATA>
34	2005-11-05/06:00:00	98	0	1	8	543	18.0	278.8	41.7	-8.8 <BUFR DATA>



The TIDB2 Browser: KTIDBExplorer

The screenshot shows the KondBExplorer application window. The title bar reads "KondBExplorer". The menu bar includes "File", "Edit", "Window", and "Help". Below the menu bar is a toolbar with various icons. The left pane shows a tree view with the following structure:

- bufr
 - local
 - /tmp/bufr2
 - bufr.new
 - /tmp/bufrs
 - grib
 - local
 - 4D01.ATL_
 - 4D01.IBER0
 - 4D01.VAM1
 - 4D01.VPNC
 - 4D02.ATL_
 - 4D02.ATL_
 - 4D02.IBER0
 - 4D02.VAM1
 - 4D02.VPNC
 - 4D02.VPN
 - 4D03.ATL_
 - 4D03.IBER0
 - 4D03.VAM1
 - 4D03.VPNC
 - 4D04.ATL_
 - 4D04.ATL_
 - 4D04.IBER0
 - 4D04.MUN
 - 4D04.VAM1
 - 4D04.VPNC
 - 4D04.VPN
 - 4D05.ASIA
 - 4D05.ATL_
 - 4D05.ATL_

The main area displays two data tables. The top table is titled "/4D01.IBER05.ECMF" and has columns: Date, Forecast, Level, Lat, Long, and W. It contains three rows of data.

	Date	Forecast	Level	Lat	Long	W
1	2005-11-09/18:00.00	2005-11-09/18:00.00	850	45.0	-10.0	0.0456
2	2005-11-09/18:00.00	2005-11-09/18:00.00	850	45.0	-9.5	0.0964
3	2005-11-09/18:00.00	2005-11-09/18:00.00	850	45.0	-9.0	0.1525

The bottom table is titled "/bufr.new" and has columns: Parameter and Value. It contains 29 rows of data.

	Parameter	Value
1	WmoBlockNumber	8
2	WmoStationNumber	509
3	TypeStation	1
4	Year	2005
5	Month	10
6	Day	17
7	Hour	0
8	Minute	0
9	LatitudeHighAccuracy	38.77
10	LongitudeHighAccuracy	-27.1
11	HeightStation	54
12	Pressure	101350
13	PressureReducedMeanSeaLevel	102010
14	3HourPressureChange	30
15	CharacteristicPressureTendency	0
16	WindDirection10M	0
17	WindSpeed10M	1
18	DryBulbTemperature2m	287.1
19	DewPointTemperature2m	282.5
20	RelativeHumidity	<missing>
21	HorizontalVisibility	10000
22	PresentWeather	2
23	PastWeather1	1
24	PastWeather2	1
25	CloudCoverTotal	10
26	VerticalSignificanceSurfaceObserv	1
27	CloudAmount	1
28	HeightBaseCloud	800
29	CloudType	38

How to get TIDB2?

- To download tidb2 from CVS:
- `cvs -d:pserver:anonymous@cvs.sourceforge.net:/cvsroot/t-i-db login`
- `cvs -z3 -d:pserver:anonymous@cvs.sourceforge.net:/cvsroot/t-i-db co -P \tidb2`
- To download ktidbexplorer from CVS:
- `cvs -z3 -d:pserver:anonymous@cvs.sourceforge.net:/cvsroot/t-i-db co -P \ktidbexplorer`
- Tarballs can be found at:
- https://sourceforge.net/project/showfiles.php?group_id=117005
- To contact me for help:
- Email to joao.simoed@meteo.pt

The End

- Thanks to ECMWF