

# RTP Format Specification and User's Guide

Version 1.01

Howard E. Motteler

July 16, 2001

## **Abstract**

We present a data format for driving radiative transfer calculations and manipulating atmospheric profiles. Calculated and observed radiances may be included as optional fields, allowing for the representation of basic co-location datasets. An implementation as HDF 4 Vdatas is given, including Fortran, C, and Matlab application interfaces.

## **1 Introduction**

The “Radiative Transfer Profile” (RTP) format is a data format for sets of atmospheric profiles, optionally paired with calculated and/or observed radiances. The format consists of a header record and an array of profile records. It was derived from the GENLN2 user profile format, extended with selected AIRS level 2 field definitions. RTP is currently implemented as HDF 4 Vdatas and as structure arrays in Fortran, C, and Matlab.

The format is intended to give a well-defined interface to radiative transfer codes, allowing for the specification of just the information needed for such calculations. It allow for modularity of both radiative transfer codes and of other tools for manipulating profiles, including tools for field selection, level interpolation and level-to-layer translations, translation of units, and building composite profiles from multiple sources. The RTP specification has some flexibility in the field set actually saved to disk, both to save space and to provide compatibility across file versions. The optional observation fields may be used to build simple co-location datasets.

## 2 The RTP format definition

The RTP format consists of a header record with information about all profiles in a file, and one or more profiles saved as an array of records. Field definitions for the header and profile records are given below. These names are both the names of the Vdata fields and the Fortran and C structure fields, with the exception of the constituent arrays, as discussed below. Depending on the application, only a subset of the fields described here need be present in an RTP file. Fields are matched by field name, and no particular order for the header or profile fields is assumed.

### 2.1 Data Types

All the RPT fields are either 32-bit integers or 32-bit floats, as noted in the field tables, with the exception of the time fields wich are 64-bit floats. These correspond to the HDF type codes DFNT\_INT32, DFNT\_FLOAT32, and DFNT\_FLOAT64, the HDF C types int32, float32, and float64, and the Fortran types integer\*4, real\*4, and real\*8.

### 2.2 Levels and Layers

The header field `ptype` flags the profile as being a level profile, a layer profile, or a profile using AIRS pseudo-layers, as follows

- |                       |             |
|-----------------------|-------------|
| 1. level profile      | LEVPRO = 0  |
| 2. layer profile      | LAYPRO = 1  |
| 3. AIRS pseudo-layers | AIRSLAY = 2 |

### RTP Header Fields

field name	short description	data type	units
ptype	profile type	scalar int32	see note [1]
pfields	profile field set	scalar int32	see note [2]
pmin	min plevs value	scalar float32	millibars
pmax	max plevs value	scalar float32	millibars
ngas	number of gases	scalar int32	[0,MAXGAS]
glist	constituent gas list	ngas int32	HITRAN gas ID
gunit	constituent gas units	ngas int32	gas unit code
nchan	number of channels	scalar int32	count
ichan	channel numbers	nchan int32	[0,MAXCHAN]
vchan	channel center freq.	nchan float32	1/cm
mwchan	number of MW channels	scalar int32	count
mwfchan	MW channel freq's	wmchan float32	GHZ
undef1	spare, user-defined	scalar float32	undefined
undef2	spare, user-defined	scalar float32	undefined

Notes:

[1] ptype values are

1. level profile           LEVPRO = 0
2. layer profile           LAYPRO = 1
3. AIRS pseudo-layers   AIRSLAY = 2

[2] RTP profile fields are organized in five groups

1. profile data           PROFBIT = 1
2. calculated IR radiances   IRCALCBIT = 2
3. observed IR radiances   IROBSVBIT = 4
4. calculated MW Tb       MWCALCBIT = 8
5. observed MW Tb       MWOBSVBIT = 16

For example, a profile with both calculated and observed IR radiances would have pfields = PROFBIT + IRCALCBIT + IROBSVBIT

RTP Profile Fields -- Profile Data

field name	short description	data type	units
-----	-----	-----	-----
plat	profile latitude	scalar float32	[-90,90] deg.
plon	profile longitude	scalar float32	[-180,360] deg.
ptime	profile time	scalar float64	TAI
stemp	surface temperature	scalar float32	Kelvins
nrho	number of refl. pts	scalar int32	[0,MAXRHO]
rho	surface reflectance	nrho float32	[0,1]
rfreq	reflectance freq's	nrho float32	1/cm
nemis	number of emis. pts	scalar int32	[0,MAXEMIS]
emis	surface emissivities	nemis float32	[0,1]
efreq	emissivity freq's	nemis float32	1/cm
salti	surface altitude	scalar float32	meters
spres	surface pressure	scalar float32	millibars
nlevs	number of press lev's	scalar int32	[0,MAXLEV]
plevs	pressure levels	nlevs float32	millibars
plays	pressure layers	nlevs float32	millibars
palts	level altitudes	nlevs float32	meters
ptemp	temperature profile	nlevs float32	Kelvins
gas_<i>	constituent amounts	nlevs float32	PPMV
co2ppm	CO2 mixing ratio	scalar float32	PPMV
scanang	sat scan angle	scalar float32	[0,90] deg.
satzen	sat zenith angle	scalar float32	[0,90] deg.
satazi	sat azimuth angle	scalar float32	[-180,360] deg.
solzen	sun zenith angle	scalar float32	[0,90] deg.
solazi	sun azimuth angle	scalar float32	[-180,360] deg.
cfrac	cloud fraction	scalar float32	[0,1]
ctype	cloud type code	scalar int32	see text
cemis	cloud top emissivity	scalar float32	[0,1]
cprtop	cloud top pressure	scalar float32	millibars
cprbot	cloud bottom pressure	scalar float32	millibars
cngwat	cloud non-gas water	scalar float32	PPMV
cpsize	cloud particle size	scalar float32	microns
landfrac	land fraction	scalar float32	[0,1]
wspeed	wind speed	scalar float32	meters/sec
wsource	wind source	scalar float32	[-180,360] deg.

RTP Profile Fields -- MW Profile Data

field name	short description	data type	units
-----	-----	-----	-----
mwnemis	number of MW emis pts	scalar int32	[0,MWMAXEMIS]
mwfreq	MW emissivity freq's	mwnemis float32	GHZ
mwemis	MW emissivities	mwnemis float32	[0,1]
mwnstb	number of MW surf pts	scalar int32	[0,MWMAXSTB]
mwsfreq	MW surface Tb freq's	mwnstb float32	GHZ
mwstb	MW surface Tbs	mwnstb float32	kelvins

RTP Profile Fields -- Radiance Data

field name	short description	data type	units
-----	-----	-----	-----
rcalc	calculated IR rad.	nchan float32	mW/m <sup>2</sup> /cm <sup>-1</sup> /str
mwcalc	calculated MW BT	mwnchan float32	kelvins
rlat	radiance obs lat.	scalar float32	[-90,90] deg.
r lon	radiance obs lon.	scalar float32	[-180,360] deg.
rtime	radiance obs time	scalar float64	TAI
robs1	observed IR rad.	nchan float32	mW/m <sup>2</sup> /cm <sup>-1</sup> /str
irinst	IR instrument code	scalar int32	instrument code
mwobs	observed MW BT	mwnchan float32	kelvins
mwinst	MW instrument code	scalar int32	instrument code
findex	file (granule) index	scalar int32	index
atrack	along-track index	scalar int32	index
xtrack	cross-track index	scalar int32	index
undef1	spare, user-defined	scalar float32	undefined
undef2	spare, user-defined	scalar float32	undefined

For level profiles, `nlevs` is the number of levels and `plevs` the pressure levels; the `nlevs` temperature and constituent fields contain level values. For layer profiles, `nlevs` is the number of layer boundaries, `plevs` is the boundary pressures, and the `nlevs-1` temperature and constituent fields contain layer values. The optional `plays` field may be used for a nominal layer pressure, if desired. The `palts` field, if used, is altitudes for the pressure levels, for either level or layer profiles.

A convention that lower indices correspond to lower pressures is suggested but not required. The header fields `pmax` and `pmin` are intended to hold the max and min level pressures over all profiles, or some upper and lower bound on these values.

## 2.3 Field Groups

The `pfields` field in the header encodes what sets of fields are in the file. This field is used by the C/Fortran API to control what which fields will be written to a file. Profile fields are organized as five groups,

- |                                   |                             |
|-----------------------------------|-----------------------------|
| 1. profile data                   | <code>PROFBIT = 1</code>    |
| 2. calculated IR radiances        | <code>IRCALCBIT = 2</code>  |
| 3. observed IR radiances          | <code>IROBSVBIT = 4</code>  |
| 4. calculated MW brightness temp. | <code>MWCALCBIT = 8</code>  |
| 5. observed MW brightness temp.   | <code>MWOBSVBIT = 16</code> |

These groups can occur in any combination. The associated parameter names and numbers are bit fields, set in `pfields` if the associated data is present in the file. Thus for example profile data plus calculated and observed IR radiances would be represented as `pfields = PROFBIT + IRCALCBIT + IROBSVBIT`, while a profile with calculated and observed MW radiances but no IR data would have `pfields = PROFBIT + MWCALCBIT + MWOBSVBIT`.

Note that we might have `nchan > 0` without having either calculated or observed radiances in a file, to specify a set of channels whose radiances are to be calculated later.

## 2.4 Constituents

Constituent fields are named with their HITRAN gas ID's, with `gas_1` water, `gas_2` CO<sub>2</sub>, and so on. A list of HITRAN gas ID's is given in an appendix.

The header field `glist` gives a list of the constituent ID's for the constituents present in the file. The default constituent unit is PPMV.

Note that only a small subset of possible constituents are typically recognized and processed by fast models for radiative transfer calculation, typically water, ozone, and perhaps methane, CO<sub>2</sub>, and CO; see the documentation of the relevant radiative transfer code for more information.

The current Fortran and C application interfaces represent constituents as a 2D array `gamnt` whose rows are layers and whose columns are gas ID index, rather than as a set of separate fields `gas_<i>` as they are actually saved in the file; the `gas_<i>` fields are the columns of the 2D `gamnt` array.

There are a wide variety of constituent units in current use; in consideration of this we have added a `gunit` array to the header, assigning a unit code for each constituent and allowing at least the potential for automatic conversions. These unit codes are given in `gas_units_code.txt`; a version is included with the RTP distribution, but the most up-to-date version can be found in the SARTA fast radiative transfer package.

## 2.5 HDF Attributes

Attributes are associated either with the header or with the profile record set, and have three parts: the field the attribute is associated with, the attribute name, and the attribute text. In addition to proper field names, the field name “header” is used for general header attributes, and “profile” for general profile attributes.

RTP attributes should typically include such information as title, author, date, and at least a brief descriptive comment. This general information should be set as attributes of the header record. Note that the Fortran/C API uses the 2D `gamnt` array for constituents; this is not actually a vdata field, and so can not take an attribute. Attributes may be attached to individual constituents with their `gas_<i>` names, where `<i>` is the HITRAN gas ID.

## 2.6 Field Sets and Sizes

Individual profiles may have varying pressures levels, emissivity, reflectance, and surface brightness temperature sets. All profiles in a file are assumed to have the same constituent set, and if radiances are present all profiles have the same channel set.

RTP fields may be scalars or one-dimensional arrays; this is a limitation of the underlying HDF vdata format. All arrays have an associated size field. If this size field is in the header, as in the case of `ngas` or `nchan` then it is assumed to be the same for all profiles, while if the size field is in a profile, as in the case of `nlevs` or `nemis`, then it applies only to that profile.

The size of array fields in the RTP HDF vdata implementation may in some cases be bigger than what is specified by the associated size field. This can happen because the HDF vdata format requires a single size be associated with each field, which then has to be at least the max of all the actual field sizes. Because of this, when a size-field is present in its value should be used instead of the possibly larger Vdata field size.

The field set for RTP is not required to be fixed to precisely the fields listed here. Fields are matched by field name, and no particular order for the header or profile fields is assumed

## 3 Application Interfaces

### 3.1 The Fortran API

The Fortran API consists of four routines: `rtpopen`, `rtpread`, `rtpwrite`, and `rtpclose`; “man page” style documentation for these is included in an appendix. The `rtpopen` routine matches fields in the HDF vdata file with field sets actually used in the Fortran structure, and builds sets of pointers for the header and profile field, for subsequent buffer copying. Fields can be added to the Fortran interface by adding them to the `rtp.h` and `rtpdefs.f` include files and recompiling the RTP libraries.

The Fortran API uses static structures whose fields (with a few exceptions noted below) are the same as the RTP fields defined above. Normally, only a subset of the Fortran structure fields will be written, with the header field `pfields` and the header size fields used to determine what actually goes into a file.

When reading data, if a file contains header or profile fields not in the Fortran structure definition, they are simply ignored. Fields that are defined in the Fortran structure but are not in a file are returned as “BAD”, or with the first element BAD, for vectors, while missing size fields are returned as zero.

Attributes are passed to and from the Fortran API thru the `RTPATTR`

structure array The records in this array have three fields: `fname`, the field name the attribute is to be associated with, `aname`, the attribute name, and `atext`, the attribute text. The header attribute field name should be either “header”, for a general attribute or comment, or a particular header field name. Similarly, the attribute profile field name should be either “profiles” or a specific profile field. Attribute strings need to be null-terminated, with `char(0)`, and the record after the last valid record in an attribute set should have `fname` set to `char(0)`. See `ftest1.f` for and `ftest2.f` examples of reading, writing, and updating attributes.

The Fortran structures differ from the `vdata` fields in two general ways. First, instead of a `gas_<i>` profile field for each constituent, the Fortran API uses a single array `gamnt(MAXLEV,MAXGAS)` to pass constituent amounts; the `gas_<i>` fields from the HDF file are the columns of this array.

The second difference is that the Fortran/C RTP header structure includes the following max size fields, which are not actually written to the `vdata` header.

<code>mlevs</code>	max number of levels	scalar int32	[0,MAXLEV]
<code>mrho</code>	max num of refl pts	scalar int32	[0,MAXRHO]
<code>memis</code>	max num of emis pts	scalar int32	[0,MAXEMIS]
<code>mwmemis</code>	max MW emis pts	scalar int32	[0,MWMAXEMIS]
<code>mwmstb</code>	max MW sTb pts	scalar int32	[0,MWMAXSTB]

On a read, these fields are set to the associated profile `vdata` field sizes. On a write, they are used to to set the size of the associated `vdata` profile fields. They can simply be set to the MAX limits, or to zero if the fields are not used; but using an actual max for the profile set, particularly for `mlevs`, can give a significant space savings.

A makefile is supplied to build the RTP API routines as a library file `librtp.a`. A Fortran demo makefile, “`Makefile.f77`” is also provided, to compile the F77 demo programs `ftest1.f` and `ftest2.f` and link them with the RTP libraries.

## 3.2 The C API

C users can use either the Fortran API or the lower level `pv*.c` family of routines which map the `vdata` buffers directly to a list of field names and pointers; see the `pv*.c` files for more information.

### **3.3 The Matlab API**

The RTP format was first implemented as a Matlab structure array, and the current HDF implementation is a fairly direct mapping between structure arrays and HDF 4 vdatas. The Matlab RTP API is implemented in two procedures, `rtpread.m` and `rtpwrite.m`; see the help info from these files for more information. The Matlab API is very flexible, mapping structure fields to vdata fields in a general way; this makes it possible to add fields, and to read and return in a structure array only the fields actually present in an HDF Vdata set.

## NAME

rtlopen -- Fortran interface to open RTP files

## SUMMARY

rtlopen() is used to open an HDF RTP ("Radiative Transfer Profile") file for reading or writing profile data. In addition, it reads or writes RTP header data and HDF header and profile attributes.

## FORTRAN PARAMETERS

data type	name	short description	direction
-----	----	-----	-----
CHARACTER *(*)	fname	RTP file name	IN
CHARACTER *(*)	mode	'c'=create, 'r'=read	IN
STRUCTURE /RTPHEAD/	head	RTP header structure	IN/OUT
STRUCTURE /RTPATTR/	hfatt	RTP header attributes	IN/OUT
STRUCTURE /RTPATTR/	pfatt	RTP profile attributes	IN/OUT
INTEGER	rchan	RTP profile channel	OUT

## VALUE RETURNED

0 if successful, -1 on errors

## INCLUDE FILES

rtpdefs.f -- Fortran header, profile, and attribute structures

## DISCUSSION

The valid open modes are 'r' to read an existing file and 'c' to create a new file.

HDF attributes are read and written in an array of RTPATTR structures, with one structure record per attribute. Attributes should be terminated with char(0), and are returned that way, for a read. The end of the attribute array is flagged with a char(0) at the beginning of the fname field.

## NAME

rtpread -- Fortran interface to read an RTP profile

## SUMMARY

rtpread reads a profile from an open RTP channel, and returns the data in the RTPPROF structure. Successive calls to rtpread return successive profiles from the file, with -1 returned on EOF.

## FORTRAN PARAMETERS

data type	name	short description	direction
-----	-----	-----	-----
INTEGER	rchan	RTP profile channel	IN
STRUCTURE /RTPPROF/	prof	RTP profile structure	OUT

## VALUE RETURNED

1 (the number of profiles read) on success , -1 on errors or EOF

## NAME

rtpwrite -- Fortran interface to write an RTP profile

## SUMMARY

rtpwrite writes an RTP profile, represented as the contents of an RTPPROF structure, to an open RTP channel. Successive calls write successive profiles.

## FORTRAN PARAMETERS

data type	name	short description	direction
-----	-----	-----	-----
INTEGER	rchan	RTP profile channel	IN
STRUCTURE /RTPPROF/	prof	RTP profile structure	IN

## VALUE RETURNED

0 on success, -1 on errors

#### NAME

rtpclose -- Fortran interface to close an RTP open channel

#### SUMMARY

rtpclose finishes up after reading or writing an RTP file, writing out any buffers and closing the HDF interface

#### FORTRAN PARAMETERS

data type	name	short description	direction
-----	-----	-----	-----
INTEGER	rchan	RTP profile channel	IN

#### VALUE RETURNED

0 on success, -1 on errors

## HITRAN Gas List

-----

### Gases from the 1998 HITRAN line database

1 = H2O (water vapor)	17 = HI
2 = CO2	18 = ClO
3 = O3 (ozone)	19 = OCS
4 = N2O	20 = H2CO
5 = CO	21 = HOCl
6 = CH4 (methane)	22 = N2 (nitrogen)
7 = O2 (oxygen)	23 = HCN
8 = NO	24 = CH3Cl
9 = SO2	25 = H2O2
10 = NO2	26 = C2H2
11 = NH3 (ammonia)	27 = C2H6
12 = HNO3	28 = PH3
13 = OH	29 = COF2
14 = HF	30 = SF6
15 = HCl	31 = H2S
16 = HBr	

### Gases represented only by cross-sections

51 = CC13F (CFC-11)	58 = C2Cl2F4 (CFC-114)
52 = CC12F2 (CFC-12)	59 = C2ClF5 (CFC-115)
53 = CC1F3 (CFC-13)	60 = CC14
54 = CF4 (CFC-14)	61 = ClONO2
55 = CHCl2F (CFC-21)	62 = N2O5
56 = CHClF2 (CFC-22)	63 = HNO4
57 = C2Cl3F3 (CFC-113)	

## NAME

rtpdump -- basic RTP dump utility

## USAGE

rtpdump [-achp] [-n k] rtpfile

## OPTIONS

- a dump attributes
- c dump RTP channel info
- h dump header structure
- p dump profile structure
- n <k> select profile <k> for channel or profile structure dumps; the first profile is 1

## BUGS

the output is from debug and error dump routines and is not fancy; the dumps are currently written to stderr